arXiv:1609.00834v1 [stat.ME] 3 Sep 2016

Functional Data Analysis by Matrix Completion*

Marie-Hélène Descary and Victor M. Panaretos

Institut de Mathématiques
Ecole Polytechnique Fédérale de Lausanne
e-mail: marie-helene.descary@epfl.ch, victor.panaretos@epfl.ch

Abstract: Functional data analyses typically proceed by smoothing, followed by functional PCA. This paradigm implicitly assumes that rough variation is due to nuisance noise. Nevertheless, relevant functional features such as time-localised or short scale fluctuations may indeed be rough relative to the global scale, but still smooth at shorter scales. These may be confounded with the global smooth components of variation by the smoothing and PCA, potentially distorting the parsimony and interpretability of the analysis. The goal of this paper is to investigate how both smooth and rough variations can be recovered on the basis of discretely observed functional data. Assuming that a functional datum arises as the sum of two uncorrelated components, one smooth and one rough, we develop identifiability conditions for the recovery of the two corresponding covariance operators. The key insight is that they should possess complementary forms of parsimony: one smooth and low rank (large scale), and the other banded and potentially high rank (small scale). Our conditions elucidate the precise interplay between rank, bandwidth, and grid resolution. Under these conditions, we show that the recovery problem is equivalent to a low rank matrix completion, and exploit this to construct estimators of the two covariances, without assuming knowledge of the true bandwidth or rank; we establish their consistency and rates of convergence, and use them to recover the smooth and rough components of each functional datum by best linear prediction. As a result, we effectively produce separate functional PCAs for smooth and rough variation. Our work bears similarities with low rank plus sparse matrix recovery, but is essentially different in that it is genuinely functional and rests on deterministic conditions.

AMS 2000 subject classifications: Primary 62M, 15A99; secondary 62M15, 60G17.
Keywords and phrases: Analyticity, banding, covariance operator, functional PCA, low rank, resolution, scale, smoothing.

Contents

1 Introduction . . . 2
2 Problem Statement . . . 4
3 Well-Posedness: Uniqueness and Identifiability . . . 5
3.1 Uniqueness of the Decomposition $\mathscr{R} = \mathscr{L} + \mathscr{B}$ . . . 5
3.2 Identifiability at Finite Resolution . . . 6
4 Estimation by Low Rank Matrix Completion . . . 7
5 Separation of Scales . . . 10
6 Asymptotic Theory . . . 11
7 Practical Implementation via Band–Deleted PCA . . . 13
8 Simulation Study . . . 14
8.1 Rank Selection . . . 15
8.2 Comparisons . . . 16
8.3 Prediction of the smooth curves . . . 20
9 Proofs of Formal Statements . . . 20
10 Appendix . . . 31
10.1 More on the Effect of Smoothing and PCA . . . 31
10.2 Analyticity and Uniqueness . . . 33
10.3 On the Scree Plot Approach for the Choice of Tuning Parameter . . . 35
10.4 Data Analysis: Application to Air Pollution Data . . . 38
10.5 Additional Simulation Results . . . 41
References . . . 44

*Research supported by a Swiss National Science Foundation grant.

1. Introduction

Functional Principal Component Analysis, the empirical version of the celebrated Karhunen-Loève expansion, is arguably the workhorse of Functional Data Analysis (Bosq [1], Ramsay and Silverman [18], Horvath and Kokoszka [9], Hsing and Eubank [10], Wang et al. [21]). It aims to construct a parsimonious yet accurate finite dimensional representation of n observable iid replicates $\{X_1, \ldots, X_n\}$ of a real-valued random function $\{X(t) : t \in [0,1]\}$ under study. The sought representation is in terms of a Fourier series built using the eigenfunctions $\{\varphi_k\}$ of the integral operator $\mathscr{R}$ with kernel $\mathrm{Cov}(X(t), X(s))$. Such a finite-dimensional representation is key in functional data analysis: not only does it serve as a basis for motivating methodology by analogy to multivariate statistics, but it constitutes the canonical means of regularization in regression, testing, and prediction, which are all ill-posed inverse problems when dealing with functional data; see Panaretos and Tavakoli [17] for an account of the genesis and evolution of functional PCA, and Wang et al. [21] for an overview of its manifold applications in functional data analysis.

Since the covariance operator $\mathscr{R}$ is unknown in practice, functional PCA must be based on its empirical counterpart (Dauxois et al. [7]; Bosq [1]),

$$\hat{\mathscr{R}}_n = \sum_{i=1}^{n} (X_i - \overline{X}) \otimes (X_i - \overline{X}), \qquad \text{where } \overline{X} = \frac{1}{n}\sum_{i=1}^{n} X_i.$$

Even this, however, is seldom accessible: one cannot perfectly observe the complete sample paths of $\{X_1, \ldots, X_n\}$. Instead, one has to make do with discrete measurements

$$X_{ij} = X_i(t_j) + \varepsilon_{ij}, \qquad i = 1, \ldots, n,\ j = 1, \ldots, K, \tag{1.1}$$

where the points t_j can be random or deterministic and the array ε_{ij} is assumed to be comprised of centred iid perturbations, independent of the X_i (see, e.g. Ramsay and Silverman [18], Hall et al. [8], Li and Hsing [13]). Roughly speaking, there are two major approaches to deal with discrete measurements: to smooth the discretely observed curves and then obtain the covariance operator and spectrum of the smooth curves; and the converse, i.e. to first obtain a smoothed estimate of the covariance operator and to use this to estimate the unobservable curves and their spectrum.

The first general approach was popularised by Ramsay and Silverman [18], by means of smoothing splines, and is widely used, chiefly when the observation grid $\{t_1, \ldots, t_K\}$ is sufficiently dense. One defines smoothed curves $\widetilde{X}_i$ as

$$\widetilde{X}_i(t) = \arg \min_{f \in C^2[0,1]} \left\{ \sum_{j=1}^{K} (f(t_j) - X_{ij})^2 + \tau \|\partial_t^2 f\|_{L^2}^2 \right\}, \qquad i = 1, \ldots, n, \tag{1.2}$$

for $C^2[0,1]$ the space of twice continuously differentiable functions on $[0,1]$, and $\tau > 0$ a regularising constant. The proxy curves $\{\widetilde{X}_i\}$ are used in lieu of the unobservable $\{X_i\}$ in order to construct a "smooth" empirical covariance operator $\widetilde{\mathscr{R}}$, and the curves $\{\widetilde{X}_i\}$ are finally projected onto the span of the first r eigenfunctions of $\widetilde{\mathscr{R}}$.

A second general approach, Principal Analysis by Conditional Expectation (PACE), was introduced by Yao et al. [22] (see also Yao et al. [23]), motivated by the need to consider situations where the grid is sparse and curves are sampled at varying grid points. In our sampling setup, and assuming the array $\{\varepsilon_{ij}\}$ to be iid of variance σ^2, they exploit the fact that the $K \times K$ covariance matrix of the vector $(X_{i1}, \ldots, X_{iK})^\top$ equals (up to a factor) $\rho(t_i, t_j) + \sigma^2 \mathbf{1}\{i = j\}$. Thus, the effect of the term ε is restricted to the addition of a σ^2-ridge to the diagonal. Yao et al. [22] then delete the diagonal $i = j$ of the empirical covariance matrix of $\{X_{ij}; i = 1, \ldots, n;\ j = 1, \ldots, K\}$ and smooth what remains to obtain a smooth estimate $\widetilde{\rho}(s,t)$ of the kernel $\rho(s,t)$. The smoothing assumes (and induces) C^2-level behaviour near $t = s$. The kernel $\widetilde{\rho}(s,t)$ is then used to construct mean-square optimal predictors $\{\widetilde{X}_1, \ldots, \widetilde{X}_n\}$ of the unobservable sample paths, truncated to belong to the span of the first r eigenfunctions of $\tilde{\rho}(s,t)$.

Proceeding in either of these two ways essentially consigns any variations of smoothness class less than C^2 to pure noise, and subsequently smears them by means of smoothing; any further rough variations are

expected to be negligible, and due to small fluctuations around eigenfunctions of order at least $r+1$ (thus orthogonal to the smooth variations) and are also discarded post-PCA.

Mathematically speaking, "smooth-then-PCA" approaches correspond to an underlying ansatz that $X(t)$ is well approximated by the sum of two uncorrelated components: a "true signal" $Y(t)$ of (essentially) finite rank r and of smoothness class C^k ($k \geq 2$) and a noise component $W(t)$ whose covariance kernel is a scaled delta function $\sigma^2\delta(s-t)$, corresponding to white noise:

$$X_i(t) = Y_i(t) + W_i(t), \qquad i = 1, \ldots, n, \tag{1.3}$$

$$X_{ij} = Y_i(t_j) + W_i(t_j) = Y_i(t_j) + \varepsilon_{ij}, \qquad i = 1, \ldots, n\,; j = 1, \ldots, K. \tag{1.4}$$

The first equation can formally be understood only in the weak sense as an SDE, and in reality W would have a covariance supported on some band $\{|t-s| < \delta\}$ for some infinitesimally small $\delta > 0$. The construction of the rank r version (by PCA) of the smoothed curves $\{\widetilde{X}_i(t)\}$ can thus be seen as an the estimation of the unobservable $\{Y_i(t)\}$. Any residual variation is then indirectly attributed to W_i, seen as functional residuals, and subsequently ignored.

It may very well happen, though, that W be rough but still be mean-square continuous, possessing a covariance kernel $b(s,t) = b(s,t)\mathbf{1}\{|t-s| < \delta\}$, for b a continuous non-constant function and $\delta > 0$ non-negligible: "*the functional variation that we choose to ignore is itself probably smooth at a finer scale of resolution*" (Ramsay and Silverman [18, Section 3.2.4]). In this case, the rough variations are not due to pure noise, but to actual signal, and contain second order structure that we may not wish to confound with that of Y or discard. Quite to the contrary, it should be fair game for functional data analysis to aim to deal with variations at smaller scales δ; to quote Ramsay and Silverman [18, Section 3.2.4] again: "*this can pay off in terms of better estimation, and this type of structure may be in itself interesting; a thoughtful application of functional data analysis will always be open to these possibilities*".

To accommodate a non-trivial kernel $b(s,t)$, the smoothing spline approach would need to replace the "uncorrelated" objective function in Equation 1.2, with the "correlated" version

$$\widetilde{X}_i(t) = \arg\min_{f \in C^2[0,1]} \left\{ (f(t_1), \ldots, f(t_K))B^{-1}(f(t_1), \ldots, f(t_K))^\top + \tau \|\partial_t^2 f\|_{L^2}^2 \right\}, \tag{1.5}$$

for B the covariance matrix of $(W_i(t_1), \ldots, W_i(t_K))^\top$. Unfortunately, B is unknown, and worse still, B and $X_i(t)$ are not jointly identifiable without further (parametric) restrictions (see Opsomer et al. [15]). Similarly, the PACE approach would need to remove a non-trivial band around the diagonal of the empirical covariance operator prior to smoothing; this would lead to unidentifiability and subsequent inconsistency without further assumptions. It would seem that the two approaches cannot be remedied by means of a simple modification, and a novel approach would be needed.

The aim of the paper is to put forward such a novel approach and to fill this gap. Without assuming knowledge of the rank r or the scale δ, we set out to:

1. Determine nonparametric conditions under which the smooth and rough variation are jointly identifiable on the basis of discrete data, and elucidate how the effective rank r of the smooth component, the scale δ of the rough component, and the grid resolution K affect identifiability.
2. Construct consistent estimators of the covariance structure of Y and W, and of their *separate* functional PCA decompositions (equivalently, separating the component in X attributable to Y from that attributable to W) on the basis of n curves sampled discretely at a grid of resolution K.

We formulate the problem rigorously in Section 2. Though it might seem that a smooth-plus-rough decomposition is neither unique nor identifiable (except under parametric conditions), we demonstrate in Section 3 that under nonparametric conditions on the covariances of Y and W, such a decomposition is indeed unique (Section 3.1, Theorem 1) and moreover identifiable on the basis of discrete measurements (Section 3.2, Theorem 2). These elucidate the interplay of rank, scale, and grid resolution. Estimators of the covariances of Y and W (without assuming knowledge of the rank r and scale δ) are then constructed in Section 4 by means of band deletion and low rank matrix completion using nonlinear least squares (combining smoothing and dimension reduction into a single step). Their consistency and convergence rates is derived in Section 6. These estimates are then used in Section 5 to recover the separate functional PCAs of the Y_i and the W_i,

producing a separation of the two scales of variation. The finite sample performance of the methodology is investigated by means of a simulation study in Section 8. Section 9 collects all the proofs of our formal results. Finally, the Appendix (Section 10) contains additional discussion, examples, theoretical results, simulations, as well as a data analysis to illustrate the methodology. Sample R and Matlab Code for the implementation of our methodology can be found at http://smat.epfl.ch/code/FDA_MatrixCompletion.zip.

2. Problem Statement

Let $X : [0,1] \to \mathbb{R}$ be a mean-zero mean square continuous random function, viewed as a random element of the space of integrable real functions defined on $[0,1]$, say $L^2([0,1])$, with the usual inner product and induced norm

$$\langle f, g \rangle_{L^2} = \int_0^1 f(t)g(t)dt \qquad \& \qquad \|f\|_{L^2}^2 = \langle f, f \rangle_{L^2}.$$

Assume that X can be decomposed as

$$X(t) = Y(t) + W(t), \qquad t \in [0,1], \tag{2.1}$$

where Y and W are *uncorrelated* random functions corresponding to a "smooth" and a "rough" component respectively. This implies an additive decomposition of X's covariance operator $\mathscr{R}$, and of its integral kernel $\rho(s,t) = \mathbb{E}[X(s)X(t)]$, as

$$\mathscr{R} = \mathscr{L} + \mathscr{B}, \tag{2.2}$$

$$\rho(s,t) = \ell(s,t) + b(s,t), \qquad s,t \in [0,1], \tag{2.3}$$

respectively, where the terms on the right are the covariance operators, and kernels, of Y and W, respectively:

$$\ell(s,t) = \mathbb{E}[Y(s)Y(t)] - \mathbb{E}[Y(s)]\mathbb{E}[Y(t)], \tag{2.4}$$

$$b(s,t) = \mathbb{E}[W(s)W(t)] - \mathbb{E}[W(s)]\mathbb{E}[W(t)]. \tag{2.5}$$

We will understand the smoothness in Y to represent smooth variation of X, i.e. large scale variation occurring over the entire $[0,1]$. On the other hand, the roughness of W corresponds to variations that occur at *scales distinctly smaller* than the global scale $[0,1]$, but not necessarily the instantaneous time scale that characterizes white noise: variation that is smooth only at *shorter time scales*.

Heuristically, if $\mathscr{B}$ is to capture variation at short time scales only, say at scales of order $\delta \in (0,1)$, we expect its kernel to vanish outside a band of size δ,

$$b(s,t) = 0, \qquad \forall\, |s-t| \geq \delta.$$

Of course, it will still admit a Mercer decomposition

$$b(s,t) = \sum_{j=1}^{\infty} \beta_j \psi_j(s)\psi_j(t) = \mathbf{1}\{|t-s| < \delta\} \sum_{j=1}^{\infty} \beta_j \psi_j(s)\psi_j(t),$$

for an orthonormal system of eigenfunctions $\{\psi_j\}$. On the other hand, since $\mathscr{L}$ captures global and smooth variation features, it cannot be allowed to have localised eigenfunctions: these should be smooth enough to be *essentially global*. At the same time, they should be finitely many, otherwise they may still succeed in spanning local variations[1]. We thus postulate that

$$\ell(s,t) = \sum_{j=1}^{r} \lambda_j \eta_j(s)\eta_j(t),$$

[1] since there exist infinitely smooth orthonormal systems that are complete in $L^2[0,1]$. To be more precise, what one needs is an exponential rate of decay of the eigenvalues $\{\lambda_j\}$, rather than a precisely finite rank, but we will see in Section 3 that a fast rate of decay alone would not suffice for identifiability to hold.

for $r < \infty$ and for $\{\eta_j\}_{j=1}^r$ sufficiently smooth orthonormal functions in $L^2[0,1]$. We will refer to the operator $\mathscr{L}$ as the *smooth operator*, and to $\mathscr{B}$ as the *banded operator*.
In summary, our setup is:

$$\rho(s,t) = \sum_{j=1}^{r} \lambda_j \eta_j(s)\eta_j(t) + \sum_{j=1}^{\infty} \beta_j \psi_j(s)\psi_j(t),$$

where: (1) $0 < \delta < 1$; (2) $r < \infty$; (3) the $\{\eta_j\}$ are sufficiently smooth. The statistical problem then is: given K discrete measurements on each of n independent copies of X,

$$X_{ij} = X_i(t_j) = Y_i(t_j) + W_i(t_j), \quad i = 1, \ldots, n,$$

obtained by point evaluation at some grid points $\{t_1, \ldots, t_K\}$,

1. estimate the components $\mathscr{L}$ and $\mathscr{B}$, and their spectral decomposition, and
2. construct separate functional PCAs for the smooth and rough components $\{Y_i\}_{i=1}^n$ and $\{W_i\}_{i=1}^n$ on the basis of these estimates (effectively separating the two scales of variation and recovering the Y_i and W_i).

To do so, we will need to formulate more precise conditions on the smoothness and roughness of the two components, or equivalently the rank and scale of these variations, as it is clear that the problem can otherwise be severely ill-posed (in a sense, the problem can be seen as an infinite-dimensional version of *density estimation with contamination by measurement error of an unknown distribution*, also known as *double-blind deconvolution*). This is done next, in Section 3.

3. Well-Posedness: Uniqueness and Identifiability

3.1. Uniqueness of the Decomposition $\mathscr{R} = \mathscr{L} + \mathscr{B}$

An obvious challenge with a decomposition of the form $\mathscr{R} = \mathscr{L} + \mathscr{B}$, is that there may be infinitely many distinct pairs $(\mathscr{L}, \mathscr{B})$ whose sum yields the same $\mathscr{R}$: we are asking to identify two summands from knowledge of their sum! As it turns out, uniqueness is a matter of scale: assuming that variations of the W process propagate only locally, at most at scale δ, whereas that variations of Y are purely non-local. The next theorem makes this statement precise via the notion of *analyticity*.

Theorem 1 (Uniqueness). *Let $\mathscr{L}_1, \mathscr{L}_2 : L^2[0,1] \to L^2[0,1]$ be trace-class covariance operators of rank $r_1 < \infty$ and $r_2 < \infty$, respectively. Let $\mathscr{B}_1, \mathscr{B}_2 : L^2[0,1] \to L^2[0,1]$ be banded trace-class covariance operators of bandwidth $\delta_1 < 1$ and $\delta_2 < 1$ respectively. If the eigenfunctions of $\mathscr{L}_1$ and $\mathscr{L}_2$ are real analytic, then we have the equivalence*

$$\mathscr{L}_1 + \mathscr{B}_1 = \mathscr{L}_2 + \mathscr{B}_2 \iff \mathscr{L}_1 = \mathscr{L}_2 \quad \& \quad \mathscr{B}_1 = \mathscr{B}_2.$$

Remark 1 (Sufficiency vs Necessity). *The conditions of the theorem can actually be strictly weakened, with the same conclusion: instead of requiring finite ranks and analytic eigenfunctions for $(\mathscr{L}_1, \mathscr{L}_2)$, it suffices to require the weaker condition that their kernels be analytic on an open set $U \subset [0,1]^2$ that contains the larger of the two bands, $U \supset \{(s,t) \in [0,1]^2 : |t-s| \leq \max(\delta_1, \delta_2)\}$. This can be relaxed no further, though: if the kernels of $(\mathscr{L}_1, \mathscr{L}_2)$ are not analytic on such a U, one can construct counterexamples, at least at this level of generality. For such counterexamples, see the Appendix (Section 10.2). Thus analyticity is necessary, unless further assumptions are imposed on the banded covariances. We choose to put the spotlight on the stronger assumption of the finite rank analytic eigenfunction case, because: (a) this is the one that will be practically relevant in light of the identifiability conditions that will be established in Section 3.2 (Theorem 2), and (b) the set of rank r covariance operators with analytic eigenfunctions is a dense subset of the set of all rank r covariance operators (see Proposition 1 below), giving us a rich set of identifiable models of the form 2.1.*

Recall that a function is real analytic on an open interval if and only if its Fourier coefficients decay at a rate that is at least geometric (see Krantz and Parks [12] for a detailed survey of real analytic functions).

For instance, if we write $\eta(x) = \sum_{k=1}^{\infty}(\alpha_k \cos(kx) + b_k \sin(kx))$, then η is real analytic on $(-\pi, \pi)$ if an only if

$$\limsup_{k\to\infty} (|\alpha_k| + |\beta_k|)^{1/k} < 1.$$

Examples of analytic functions include polynomials, trigonometric functions, exponential and logarithmic functions, rational functions with no poles, truncated Gaussians and finite location/scale mixtures thereof, to name only a few; such functions have been routinely used as typical examples of low order eigenfunctions capturing smooth variation in functional data analysis. The class of real analytic functions is also closed under finite linear combination, multiplication and division (assuming a non-vanishing denominator), composition, differentiation, and integration. Thus, one can generate rich collections of analytic eigenfunctions (and hence analytic covariance operators) by combining analytic functions. In fact, the set of rank r covariance operators with analytic eigenfunctions is a dense subset of the set of all rank r covariance operators:

Proposition 1. *Let Z be an $L^2[0,1]$-valued random function with a trace class covariance $\mathscr{G}$ of rank $r < \infty$. Then, for any $\epsilon > 0$ there exists a random function Y whose covariance $\mathscr{L}$ has analytic eigenfunctions and rank $q \leq r$, such that*

$$\mathbb{E}\|Z - Y\|_{L^2}^2 < \epsilon, \quad \& \quad \|\mathscr{G} - \mathscr{L}\|_* < \epsilon,$$

for $\|\cdot\|_*$ the nuclear norm. If additionally $\mathscr{G}$ has C^1 eigenfunctions on $[0,1]$, then we have the stronger result that for any $\epsilon > 0$, there exists a random function Y whose covariance $\mathscr{L}$ has analytic eigenfunctions and rank $q \leq r$, such that

$$\sup_{t\in[0,1]} \mathbb{E}|Z(t) - Y(t)|^2 < \epsilon, \quad \& \quad \sup_{s,t\in[0,1]} |g(s,t) - \ell(s,t)| < \epsilon,$$

where g and ℓ are the kernels of $\mathscr{G}$ and $\mathscr{L}$, respectively.

Note that an immediate conclusion is that, for a given r, the accuracy of a rank r analytic approximation of a mean-square continuous process can be made arbitrarily close to the accuracy of the (optimal) rank r Karhunen-Loève approximation, in the same uniform mean square sense. Thus, if we expect a process to be approximately of low rank r (as in our Model of Section 2), then this process can be very well approximated by an analytic process of the same rank low rank r. This shows that the condition of analyticity, at least as a model that guarantees uniqueness of decomposition $\mathscr{R} = \mathscr{L} + \mathscr{B}$, is not nearly as restrictive as it may seem at first sight (and in any case, it is sharp given the discussion in Remark 1).

3.2. Identifiability at Finite Resolution

Theorem 1 relies on an analyticity assumption, which is a fundamentally functional assumption, so it is not clear whether the result is useful in practice: is the decomposition identifiable on the basis of finitely many discrete measurements? Remarkably the answer is yes, and crucially depends both on the finite rank and the analyticity assumption.

Suppose we are given K discrete measurements on each of n independent copies of X,

$$X_{ij} = X_i(t_j) = Y_i(t_j) + W_i(t_j), \quad i = 1, \ldots, n,$$

obtained by evaluation at points $\{t_j\}_{j=1}^K$, where

$$(t_1, \ldots, t_K) \in \mathcal{T}_K = \left\{(x_1, \ldots, x_K) \in \mathbb{R}^K : x_1 \in I_{1,K}, \ldots, x_K \in I_{K,K}\right\},$$

and $\{I_{j,K}\}_{j=1}^K$ is the partition of $[0,1]$ into intervals of length $1/K$. With this information, we can of course only hope to be able to uniquely identify the K-resolution versions of the operators, $(\mathscr{L}, \mathscr{B})$, say $(\mathscr{L}^K, \mathscr{B}^K)$ on the basis of the K-resolution version of their sum, say $\mathscr{R}^K = \mathscr{L}^K + \mathscr{B}^K$. These operators are defined to have kernels:

$$\rho^K(x,y) \quad = \quad \sum_{i,j=1}^{K} \rho(t_i, t_j)\mathbf{1}\{(x,y) \in I_{i,K} \times I_{j,K}\}, \tag{3.1}$$

$$\ell^K(x,y) = \sum_{i,j=1}^{K} \ell(t_i,t_j)\mathbf{1}\{(x,y)\in I_{i,K}\times I_{j,K}\}, \tag{3.2}$$

$$b^K(x,y) = \sum_{i,j=1}^{K} b(t_i,t_j)\mathbf{1}\{(x,y)\in I_{i,K}\times I_{j,K}\}, \tag{3.3}$$

which can be summarised via the following $K\times K$ matrix representations:

$$R^K(i,j)=\rho(t_i,t_j),\quad L^K(i,j)=\ell(t_i,t_j),\quad B^K(i,j)=b(t_i,t_j).$$

As it turns out, there exists a finite critical resolution K^*, with explicit dependence on the rank r and scale δ, beyond which identification is possible, provided that $r<\infty$ and $\delta<1/2$. This encapsulates the interplay between rank, resolution, and scale:

Theorem 2 (Discrete Identifiability)**.** *Let $\mathscr{L}_1$ and $\mathscr{L}_2$ be covariance operators of finite ranks $r_1<\infty$ and $r_2<\infty$, respectively, and assume without loss of generality that $r_1\geq r_2$. Let $\mathscr{B}_1$ and $\mathscr{B}_1$ be two banded continuous covariance operators of bandwidth $\delta_1<1/2$ and $\delta_2<1/2$ respectively. Given $(t_1,\ldots,t_K)\in\mathcal{T}_K$, define their K-resolution matrix coefficients to be $(L_1^K,B_1^K,L_2^K,B_2^K)\in\mathbb{R}^{K\times K}$,*

$$L_m^K(i,j)=\ell_m(t_i,t_j)\quad\&\quad B_m^K(i,j)=b_m(t_i,t_j),\qquad i,j\in\{1,\ldots,K\},$$

for $m=1,2$. If the eigenfunctions of $\mathscr{L}_1$ and $\mathscr{L}_2$ are all real analytic, and

$$K\geq K^*=\max\left(\frac{2r_1+2}{1-2\delta_1},\frac{2r_1+2}{1-2\delta_2}\right),$$

then we have the equivalence

$$L_1^K+B_2^K=L_2^K+B_2^K \iff L_1^K=L_2^K\,\&\,B_1^K=B_2^K,$$

almost everywhere on $\mathcal{T}_K$ with respect to Lebesgue measure.

The theorem reveals the interplay between the fundamental parameters of the problem, which is governed by the constraint:

$$r\leq\left(\frac{1}{2}-\delta\right)K-1. \tag{3.4}$$

This yields the maximal rank that the smooth operator can have, for a given resolution K and scale δ of the banded operator, if the problem is to be identifiable. Figure 1 plots this maximal rank r as a function of K for different values of the parameter δ. We note that things are not particularly restrictive, allowing identifiability even at "sparse" resolutions K, depending of course on the value of δ. To some extent, this bridges the results in Yao et al, [22], who work in the limiting case $\delta=0$.

An attractive feature of this result is that the conditions imposed are deterministic and yet not particularly restrictive. This is in contrast with results in recent progress on matrix completion which either have restrictive deterministic conditions, or more relaxed but random conditions. The reason is that we are fortunate to have a deterministic and known structure of the missing set of values to be completed.

The main caveat of passing from the continuum to discrete observation, is that the theorem is valid almost everywhere on $\mathcal{T}_K$, rather than pointwise on $\mathcal{T}_K$. Thus, we know that the identifiability holds for *almost all grids* without being able to conclusively say so for *a specific grid*. In probabilistic terms, if the points t_j are chosen independently at random, each according to an absolutely continuous distribution on the corresponding interval I_j, then we know that identifiability holds with probability 1.

4. Estimation by Low Rank Matrix Completion

Our strategy for estimation will be to define an objective function depending only on R^K whose unique optimum yields the required matrix L^K. Then, we will define an estimator of L^K on the basis of an empirical version of this objective function. Ideally, the objective function should not depend on the knowledge of the unknown quantities δ and r, otherwise there would be two "competing" tuning parameters to choose. The following proposition yields such an objective function, in the form of a low rank matrix completion problem:

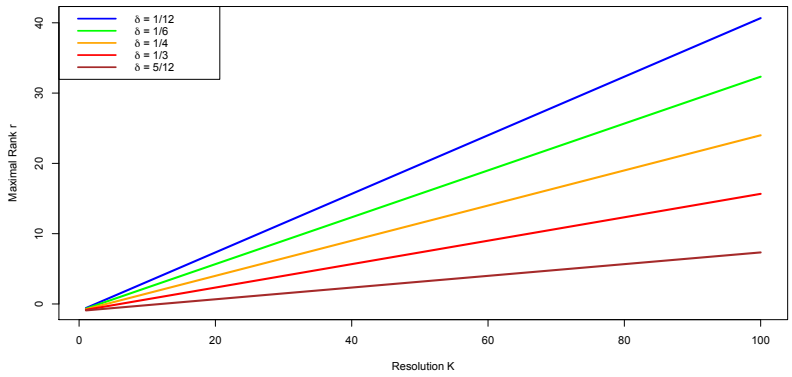

FIG 1. *Graphic representation of the interplay between rank, scale, and resolution. For different values of the scale parameter δ, the maximal identifiable rank r is plotted as a function of the resolution K.*

Proposition 2. *Let $\mathscr{L} : L^2[0,1] \to L^2[0,1]$ be a rank $r < \infty$ covariance operator with analytic eigenfunctions and kernel ℓ, and $\mathscr{B} : L^2[0,1] \to L^2[0,1]$ a trace-class covariance operator with δ-banded kernel b. For $(t_1, \ldots, t_K) \in \mathcal{T}_K$, let*

$$L^K = \{\ell(t_i, t_j)\}_{ij}, \quad B^K = \{b(t_i, t_j)\}_{ij},$$

and $R^K = L^K + B^K$. Assume that

$$\delta < \frac{1}{4} \qquad \& \qquad K \geq 4r + 4.$$

Define the matrix $P^K \in \mathbb{R}^{K \times K}$ by $P^K(i,j) = \mathbf{1}\{|i-j| > \lceil K/4 \rceil\}$. Then, for almost all grids in $\mathcal{T}_K$:

1. *The matrix L^K is the unique solution to the optimisation problem*

$$\min_{\theta \in \mathbb{R}^{K \times K}} \operatorname{rank}\{\theta\} \qquad \text{subject to } \left\| P^K \circ (R^K - \theta) \right\|_F^2 = 0. \tag{4.1}$$

2. *Equivalently, in penalised form,*

$$L^K = \operatorname*{arg\,min}_{\theta \in \mathbb{R}^{K \times K}} \left\{ \left\| P^K \circ (R^K - \theta) \right\|_F^2 + \tau \operatorname{rank}(\theta) \right\}, \tag{4.2}$$

   for all $\tau > 0$ sufficiently small.

Here $\|\cdot\|_F$ is the Frobenius matrix norm and "$\circ$" denotes the Hadamard product.

Simply put, among all possible matrix completions of $P^K \circ (R^K - \theta)$, the matrix L^K is uniquely the one of lowest rank: no matrix of rank lower than the true rank r will provide a completion; and any completion other than L^K will have rank at least $r+1$. Notice that neither of the objective functions 4.1 or 4.2 depends on δ or r: unique recovery of L^K and B^K is feasible even when we don't know the true values of r or δ. The concession we had to make to achieve this adaptation is to require $\delta < 1/4$ (compared to $\delta < 1/2$ in Theorem 2). In particular, we use the penalised form in Equation 4.2 to motivate the formal definition of our estimation approach (the equivalent form in Equation 4.1 will come in useful for computation, see Section 7):

Definition 1 (Estimator of L^K)**.** *Let $(X_1, \ldots, X_n)$ be iid copies of $X = Y + W$. Let $(t_1, \ldots, t_K) \in \mathcal{T}_K$ and assume we observe*

$$X_{ij} = X_i(t_j), \quad i = 1, \ldots, n;\ j = 1, \ldots, K.$$

Let $R_n^K \in \mathbb{R}^{K \times K}$ be the empirical covariance matrix of the vectors

$$\{(X_{i1}, \ldots, X_{iK})^\top\}_{i=1}^n.$$

We define the estimator $\hat{L}_n^K$ of L^K to be an approximate minimum of

$$\min_{0 \preceq \theta \in \mathbb{R}^{K \times K}} \quad \frac{1}{K^2} \left\{ \left\| P^K \circ (R_n^K - \theta) \right\|_F^2 + \tau_n \mathrm{rank}(\theta) \right\}, \tag{4.3}$$

where $P^K \in \mathbb{R}^{K \times K}$ is defined as $P^K(i,j) = \mathbf{1}\{|i-j| > \lceil K/4 \rceil\}$ and $\tau_n > 0$ is a tuning parameter such that $\tau_n \to 0$ as $n \to \infty$. By approximate minimum it is meant that the value of the functional at $\hat{L}_n^K$ is within $O_{\mathbb{P}}(n^{-1})$ of the value of the overall minimum.

We discuss the practical implementation of the estimation method of Definition 1, including the selection of the tuning parameter, in Section 7. While it can happen that the optimisation problem not have a unique solution at a finite sample size n, Section 6 shows that any sequence of such estimators indexed by (n, K) will be consistent, provided the grid becomes dense and the number of curves diverges. The same section also discusses the required decay rate of τ_n for consistency (see Remark 2). Once $\hat{L}_n^K$ has been constructed, we may also construct a plug-in estimator for B^K:

Definition 2 (Plug-in Estimator of B^K)**.** *Let R_n^K and $\hat{L}_n^K$ be as in Definition 1. We define the plug-in estimator $\hat{B}_n^K$ of B_n^K to be the alternated projection of $\Delta_n^K = R_n^K - \hat{L}_n^K$ onto the convex set of non-negative banded $K \times K$ matrices of bandwidth at most $\lceil K/4 \rceil$.*

We could of course have used $\Delta_n^K = R_n^K - \hat{L}_n^K$ itself to estimate B^K, but there is no guarantee that this will be positive definite. Asymptotically in n, Δ_n^K and $\hat{B}_n^K$ will coincide. Note that the intersection of the set of banded matrices (with given band) and the set of non-negative matrices is a closed convex set, thus the projection uniquely exists. In practice it can be approximately determined by the method of alternative projections, or Dykstra's algorithm (see Section 7).

Once $\hat{L}_n^K$ and $\hat{B}_n^K$ are at hand, it is reasonable to use their sum as an estimator of R^K, instead of the empirical version R_n^K, as the former is in principle less "noisy" than the latter:

Definition 3 (Plug-in Estimator of R^K)**.** *Let $\hat{L}_n^K$ and $\hat{B}_n^K$ be as in Definitions 1 and 2. We define the plug-in estimator $\hat{R}_n^K$ of R^K as $\hat{R}_n^K = \hat{L}_n^K + \hat{B}_n^K$.*

Our K-resolution estimators $(\hat{\mathscr{L}}_n^K, \hat{\mathscr{B}}_n^K, \hat{\mathscr{R}}_n^K)$ of $(\mathscr{L}, \mathscr{B}, \mathscr{R})$ will now be defined as the operators with step-function kernels $(\hat{\ell}_n^K(x,y), \hat{b}_n^K(x,y), \hat{\rho}_n^K(x,y))$ whose coefficients are given by the matrices $(\hat{L}_n^K, \hat{B}_n^K, \hat{R}_n^K)$:

$$\begin{aligned}
\hat{\ell}_n^K(x,y) &= \sum_{j=1}^K \hat{L}_n^K(i,j) \mathbf{1}\{(x,y) \in I_{i,K} \times I_{j,K}\}, \\
\hat{b}_n^K(x,y) &= \sum_{j=1}^K \hat{B}_n^K(i,j) \mathbf{1}\{(x,y) \in I_{i,K} \times I_{j,K}\}, \\
\hat{\rho}_n^K(x,y) &= \sum_{j=1}^K \hat{R}_n^K(i,j) \mathbf{1}\{(x,y) \in I_{i,K} \times I_{j,K}\}.
\end{aligned}$$

Correspondingly, the estimators of their spectra will be given by the spectra of $\hat{\mathscr{L}}_n^K$, $\hat{\mathscr{B}}_n^K$, and $\hat{\mathscr{R}}_n^K$:

$$\hat{\mathscr{L}}_n^K = \sum_{j=1}^{\hat{r}} \hat{\lambda}_j \hat{\eta}_j \otimes \hat{\eta}_j, \quad \hat{\mathscr{B}}_n^K = \sum_{j=1}^K \hat{\beta}_j \hat{\psi}_j \otimes \hat{\psi}_j, \quad \hat{\mathscr{R}}_n^K = \sum_{j=1}^K \hat{\theta}_j \hat{\varphi}_j \otimes \hat{\varphi}_j.$$

Here, $\hat{r} \leq K/4$ is the rank of $\hat{\mathscr{L}}_n^K$. Note that the empirical eigenfunctions $\hat{\eta}_j$ of $\hat{\mathscr{L}}_n^K$ will be step functions. They can, of course, be replaced by smooth versions thereof. For example, one can smooth the covariance function $\hat{\ell}_n^K$, and then calculate the spectrum of the induced covariance operator. The amount of smoothing required will be rather limited since $\hat{\ell}_n^K$ is effectively already de-noised. One could also directly smooth the eigenfunctions, but then there is no guarantee that their smoothed versions will be still orthogonal. Without any additional smoothness assumptions on $\mathscr{B}$, we cannot presume to smooth the step functions $\hat{\psi}_j$ in order to obtain smoother versions (recall that the only continuity of $\mathscr{B}$ was assumed).

5. Separation of Scales

With estimators of the covariance operators $(\mathscr{L}, \mathscr{B})$ and their spectra at our disposal, we now wish to carry out functional PCA separately for the smooth and the rough components, thus separating the two scales of variation. In order to have identifiability at the level of curves, we need to add the assumption that at least one of the two processes Y and W has a known mean. Here we assume that the rough process W is known to have mean zero, and to simplify the presentation we assume that the mean of Y has been removed from the data so we have $\mathbb{E}[Y] = 0$ too. Focussing on the smooth component, we note that its Karhunen-Loève expansion is

$$Y_i = \sum_{j=1}^{r} \langle Y_i, \eta_j \rangle \eta_j.$$

Having estimated η_j already, it suffices to estimate the scores $\{\langle Y_i, \eta_j \rangle\}_{i=1}^n$, in order to have a complete analysis into principal components. If we were able to observe $\{Y_i(t_j)\}_{i,j}$, then the natural estimator would be given by

$$\langle Y_i^K, \hat{\eta}_j \rangle_{L^2} = \frac{1}{K} \sum_{k=1}^{K} Y_i(t_k) \hat{\eta}_j(t_k),$$

where $Y_i^K(t) = \sum_{j=1}^K Y_i(t_j) \mathbf{1}\{t \in I_{j,K}\}$. A parallel discussion holds in the case of the rough components $\{W_i\}$. In effect, we see that the problem of estimating the principal scores of Y and W separately is equivalent to that of *separating* the unobservable components $Y_i(t_j)$ and $W_i(t_j)$ in the decomposition

$$X_i(t_j) = Y_i(t_j) + W_i(t_j),$$

on the basis of the observations $X_i(t_j)$. We concentrate on a specific observation, say $i = 1$, and drop the index 1 for the sake of tidiness.

Separation can be viewed as a problem of *prediction* (similarly to the approach taken by Yao et al. [22]). If the covariance operators $\mathscr{R}$ and $\mathscr{L}$ were known precisely, then we would attempt to recover the components $Y^K(t) = \sum_{j=1}^K Y(t_j) \mathbf{1}\{t \in I_{j,K}\}$ and $W^K(t) = \sum_{j=1}^K W(t_j) \mathbf{1}\{t \in I_{j,K}\}$ by means of their best predictors given the observation $X^K(t) = \sum_{j=1}^K X(t_j) \mathbf{1}\{t \in I_{j,K}\}$. The most tractable case is that of using the best *linear* predictor (which is best overall in the Gaussian case), and this is what we will pursue. Noting that Y and W are zero mean and uncorrelated, the best linear predictor of Y^K given X^K (viewed as random elements of L^2) is

$$\Pi(X^K) = \sum_{j=1}^{r} \sum_{i=1}^{q} \frac{\lambda_j^K}{\theta_i^K} \langle \varphi_i^K, \eta_j^K \rangle \langle \varphi_i^K, X^K \rangle \eta_j^K = \sum_{j=1}^{r} \xi_j \eta_j^K, \tag{5.1}$$

where $\{\theta_i^K, \varphi_i^K\}_{i=1}^q$ is the spectrum of $\mathscr{R}^K$ (with $q \leq \infty$) and $\{\lambda_j^K, \eta_j^K\}_{j=1}^r$ that of $\mathscr{L}^K$ (see Bosq [2, Prop. 3.1], and Bosq [2, Example 3.3]). Note that $\mathscr{R}^K$ is the covariance operator of X^K.

We estimate the best linear predictor, by replacing the unknown elements in Equation 5.1 by their corresponding estimators. Specifically, recalling that

$$\hat{\mathscr{R}}_n^K = \sum_{i=1}^{\hat{q}} \hat{\theta}_i \hat{\varphi}_i \otimes \hat{\varphi}_i, \quad \hat{q} = \operatorname{rank}(\hat{\mathscr{R}}_n^K) \qquad \& \qquad \hat{\mathscr{L}}_n^K = \sum_{j=1}^{\hat{r}} \hat{\lambda}_j \hat{\eta}_j \otimes \hat{\eta}_j, \quad \hat{r} = \operatorname{rank}(\hat{\mathscr{L}}_n^K),$$

our estimator of the predictor of Y^K given X^K is

$$\hat{Y}_n^K := \sum_{j=1}^{\hat{r}} \sum_{i=1}^{\hat{q}} \frac{\hat{\lambda}_j}{\hat{\theta}_i} \langle \hat{\varphi}_i, \hat{\eta}_j \rangle \langle \hat{\varphi}_i, X^K \rangle \hat{\eta}_j = \sum_{j=1}^{\hat{r}} \hat{\xi}_j \hat{\eta}_j. \tag{5.2}$$

In matrix notation, the estimated scores $(\hat{\xi}_1, \ldots, \hat{\xi}_{\hat{r}})^\top$ of Y satisfy

$$\hat{\xi}_j = \langle \hat{\lambda}_j (\hat{\mathscr{R}}_n^K)^\dagger \hat{\eta}_j, X^K \rangle = \frac{1}{K} \hat{\lambda}_j \boldsymbol{X}^\top (\hat{R}_n^K)^\dagger \hat{\boldsymbol{\eta}}_j = \frac{1}{K} \hat{\lambda}_j \boldsymbol{X}^\top (\hat{L}_n^K + \hat{B}_n^K)^\dagger \hat{\boldsymbol{\eta}}_j, \tag{5.3}$$

where $\boldsymbol{X} = (X(t_1), \ldots, X(t_K))^\top$, $\hat{\boldsymbol{\eta}}_j = (\hat{\eta}_j(t_1), \ldots, \hat{\eta}_j(t_K))^\top$, and we use the notation $\mathscr{A}^\dagger$ to denote the generalised inverse of an operator (or matrix) $\mathscr{A}$. It is worth remarking that the last expression in Equation 5.3 is essentially the same as that of the PACE estimator of Yao et al. [22], with the exception that one has a banded matrix $\hat{B}_n^K$ in lieu of a diagonal matrix of the form $\hat{\sigma}^2 I$. The best linear predictor of W^K given X^K, say $\Psi(X^K)$, can be estimated by means of the *residuals*

$$\hat{W}(t_j) = X(t_j) - \hat{Y}_n^K(t_j), \qquad j = 1, \ldots, K.$$

This definition is motivated from the simple fact that

$$\Psi(X^K) = \mathbb{E}\big[W^K \big| X^K\big] = \mathbb{E}\big[X^K - Y^K \big| X^K\big] = X^K - \mathbb{E}\big[Y^K \big| X^K\big] = X^K - \Pi(X^K).$$

6. Asymptotic Theory

We now turn to establishing the consistency of the estimators constructed in the last two Sections. Our first result derives the consistency and rate of convergence of our estimator $\hat{\mathscr{L}}_n^K$ and its spectrum, as the observation grid becomes dense and the number of curves diverges. To simplify the theoretical analysis, we will consider nested sequences of grids, obtained by refinements of the grid of *critical resolution*. In the sequel, we will follow the usual convention that the sign of the estimated eigenfunctions is correctly identified (since only the eigenprojectors are formally identifiable).

Theorem 3 (Consistency of $\hat{\mathscr{L}}$, $\hat{\eta}$, and $\hat{\lambda}$)**.** *In the setting of Section 4, let $\mathbb{E}\|X\|_{L^2}^4 < \infty$, and $\delta < \frac{1}{4}$, and define $K^* = 4(r+1)$ to be the critical resolution. Provided that $\tau_n \to 0$, it holds almost everywhere on $\mathcal{T}_K$ that*

$$\left\|\hat{\mathscr{L}}_n^K - \mathscr{L}\right\|_{\mathrm{HS}}^2 \leq O_{\mathbb{P}}(n^{-1}) + 4K^{-2} \sup_{x,y\in[0,1]} \|\nabla \ell(x,y)\|_2^2, \tag{6.1}$$

$$\|\hat{\eta}_j - \eta_j\|_{L^2}^2 \leq O_{\mathbb{P}}(n^{-1}) + 2K^{-1}\|\eta_j'\|_\infty^2, \quad j \in \{1, \ldots, r\}, \tag{6.2}$$

$$\sup_{j\geq 1} |\hat{\lambda}_j - \lambda_j|^2 = O_{\mathbb{P}}(n^{-1}) + 4K^{-2} \sup_{x,y\in[0,1]} \|\nabla \ell(x,y)\|_2^2, \tag{6.3}$$

for any refinement $K = m \times K^*$, $m \geq 1$, where $\|\cdot\|_{\mathrm{HS}}$ is the Hilbert-Schmidt norm of an operator. All three $O_{\mathbb{P}}(n^{-1})$ terms are uniform in K. Furthermore, the rank of $\hat{\mathscr{L}}_n^K$ satisfies,

$$n\tau_n K^{-2} |\mathrm{rank}(\hat{\mathscr{L}}_n^K) - r| = O_{\mathbb{P}}(1). \tag{6.4}$$

Remark 2. *If K is at least of the order $\sqrt{n}$, then we attain a parametric convergence rate for $\hat{\mathscr{L}}_n^K$, irrespective of the convergence rate of τ_n to zero. On the other hand, Equation 6.4 reveals that if we wish to consistently estimate the rank, we must fix a sufficiently large but finite grid K, and choose τ_n such that $n\tau_n \to \infty$ (or choose $K = o(1/\sqrt{n\tau_n})$ with $n\tau_n \to \infty$). It thus seems that accurate estimation of the operator requires dense grids, but accurate estimation of the rank requires sparse grids.*

Remark 3. *The fact that the theorem holds true almost everywhere on $\mathcal{T}_K$ can equivalently be stated in probabilistic terms. Assume that the grid $\boldsymbol{t}_K = \{t_{j,K}\}_{j=1}^K$ is chosen at random according to the uniform distribution on $\mathcal{T}_K$. Then the theorem holds with probability 1 over the grid choice. Note that the uniform measure on $\mathcal{T}_K$ can be generated by selecting $\{t_{j,K}\}_{j=1}^K$ to be independent for $j \in \{1, \ldots, K\}$, each uniformly distributed on the corresponding subinterval $I_{j,K}$.*

The consistency of $\hat{\mathscr{R}}_n^K$ will follow as a corollary, since it is defined as a contraction of the difference $\mathscr{R}_n^K - \hat{\mathscr{L}}_n^K$. Unless we assume some regularity of b, however, we will not be able to obtain convergence rates. If we are able to assume C^1 behaviour of b on the band $\{|x-y| < \delta\}$, we do obtain rates of convergence, which improve if C^1 behaviour can be assumed everywhere on $[0,1]^2$.

Corollary 1 (Consistency of $\hat{\mathscr{B}}$ and $\hat{\psi}$)**.** *Under the same conditions as in Theorem 3, and for any $\epsilon, \gamma > 0$ there exists $M_1(\epsilon, \gamma) > 0$ such that*

$$\mathbb{P}\left\{\left\|\hat{\mathscr{B}}_n^K - \mathscr{B}\right\|_{\mathrm{HS}} > \epsilon\right\} < \gamma, \quad \& \quad \mathbb{P}\left\{\sup_{j\geq 1}|\hat{\beta}_j - \beta_j| > \epsilon\right\} < \gamma \qquad \forall\, K, n > M_1(\epsilon, \gamma),$$

and $M_2(\epsilon, \gamma, j) > 0$ such that

$$\mathbb{P}\left\{\left\|\hat{\psi}_j - \psi_j\right\|_{L^2} > \epsilon\right\} < \gamma, \qquad \forall\, K, n > M_2(\epsilon, \gamma, j),$$

for almost every grid $(t_1, \ldots, t_K) \in \mathcal{T}_K$. If the kernel $b(s,t) : [0,1]^2 \to \mathbb{R}$ of $\mathscr{B}$ is additionally assumed to be continuously differentiable on $\{|t-s| < \delta\}$, we obtain the rates of convergence

$$\left\|\hat{\mathscr{B}}_n^K - \mathscr{B}\right\|_{\mathrm{HS}}^2 = O_{\mathbb{P}}(n^{-1}) + O(K^{-1}), \tag{6.5}$$

$$\left\|\hat{\psi}_j - \psi_j\right\|_{L^2}^2 = \frac{O_{\mathbb{P}}(n^{-1}) + O(K^{-1})}{\sigma_j^2}, \tag{6.6}$$

$$\sup_{j\geq 1}|\hat{\beta}_j - \beta_j| = O_{\mathbb{P}}(n^{-1}) + O(K^{-1}), \tag{6.7}$$

for any refinement $K = m \times K^*$, $m \geq 1$, and almost every grid $(t_1, \ldots, t_K) \in \mathcal{T}_K$, where

$$\sigma_1 = \beta_1 - \beta_2, \quad \& \quad \sigma_j = \min\{\beta_{j-1} - \beta_j, \beta_j - \beta_{j+1}\}, \qquad 2 \leq j \leq \operatorname{rank}(\mathscr{B}),$$

and all three $O_{\mathbb{P}}(n^{-1})$ terms are uniform in K. If b is continuously differentiable on the whole unit square $[0,1]^2$, then the $O(K^{-1})$ terms in Equations 6.5, 6.6, and 6.7 can be improved to $O(K^{-3})$.

Remark 4. *The fact that the "bias" term can be improved $O(K^{-3})$ when $b \in C^1([0,1]^2)$ follows from the fact that B^K is banded, and thus the number of its non-zero entries grow like K rather than like K^2.*

The convergence rates for $\hat{\mathscr{L}}_n^K$ and $\hat{\mathscr{B}}_n^K$ can now be combined in order to yield convergence rates for $\hat{\mathscr{R}}_n^K = \hat{\mathscr{L}}_n^K + \hat{\mathscr{B}}_n^K$.

Corollary 2 (Convergence rates for $\hat{\mathscr{R}} = \hat{\mathscr{L}} + \hat{\mathscr{B}}$)**.** *Under the same conditions as in Theorem 3, and if the kernel $b(s,t) : [0,1]^2 \to \mathbb{R}$ of $\mathscr{B}$ is assumed to be continuously differentiable on $\{|t-s| < \delta\}$, we obtain the rate of convergence*

$$\left\|\hat{\mathscr{R}}_n^K - \mathscr{R}\right\|_{\mathrm{HS}}^2 = O_{\mathbb{P}}(n^{-1}) + O(K^{-1}), \tag{6.8}$$

for any refinement $K = m \times K^*$, $m \geq 1$, and almost all grids in $\mathcal{T}_K$. The $O_{\mathbb{P}}(n^{-1})$ term is uniform in K. If b is continuously differentiable on the whole unit square $[0,1]^2$, then the $O(K^{-1})$ term in Equation 6.8 can be improved to $O(K^{-2})$.

Finally, we show that the predictors of Y^K and W^K based on a finite grid of resolution K are consistent in the L^2 sense, which also implies that the estimated PCA scores are consistent too:

Corollary 3 (Consistency of K-Resolution Predictors)**.** *In the same setting as in Theorem 3, let $K = m \times K^*$ $(m \geq 1)$ be a refinement of the critical grid size. If $\mathscr{R}^K$, is of full rank, and if the kernel $b(s,t) : [0,1]^2 \to \mathbb{R}$ of $\mathscr{B}$ is continuously differentiable on $\{|t-s| < \delta\}$,then*

$$\|\hat{Y}_n^K - \Pi(X^K)\|_{L^2} = O_{\mathbb{P}}(n^{-1/2}),$$
$$\|\hat{W}_n^K - \Psi(X^K)\|_{L^2} = O_{\mathbb{P}}(n^{-1/2}),$$

almost everywhere on $\mathcal{T}_K$.

7. Practical Implementation via Band–Deleted PCA

We now describe how we can compute the estimators $\hat{L}_n^K$ and $\hat{B}_n^K$ in practice. Starting with $\hat{L}_n^K$:

(A) The first step is to solve the optimisation problem

$$\min_{0 \preceq \theta \in \mathbb{R}^{K \times K}} \quad \left\| P^K \circ (R_n^K - \theta) \right\|_F^2 \qquad \text{subject to} \quad \operatorname{rank}(\theta) \leq i, \tag{7.1}$$

for all i ranging in $\{1, \ldots, K/4 - 1\}$. We discuss the details of how this is done later in this section.

(B) The second step is to collect the minima of 7.1 for each rank i, the *fits* $\{f(i) : i = 1, \ldots, K/4 - 1\}$, and calculate the quantities

$$f(i) + \tau i,$$

for some choice of the tuning parameter τ. We then determine the i that minimises the above quantity, and declare the corresponding optimising matrix to be the estimator $\hat{L}_n^K$. Notice that τ being positive precludes us from overfitting by choosing a matrix of arbitrarily large rank.

A natural question is: *how does one choose the precise τ in Step (B)*? The answer is that, any choice of τ implies a choice of rank i_τ (this being the rank of the optimum corresponding to τ), and thus a fit value $f(i_\tau)$. Thus one can use the the *scree-plot* $i \mapsto f(i)$ as a guide to implicitly choose τ, by replacing step (B) with:

(B') Plot the non-increasing function $i \mapsto f(i)$, and choose a value of i to be the smallest one such that $f(i) < c$, for some threshold value c. Then declare the corresponding optimising matrix to be the estimator $\hat{L}_n^K$. Again, c being positive precludes us from overfitting by choosing an arbitrarily large rank.

Remark 5. *The solution of (B') for a certain choice of $c > 0$ is equivalent to the solution of (B) for a certain corresponding choice of τ (when the scree plot has a convex shape, as has been the case in all the simulations we carried out, there is an explicit relationship between c and τ; see the Appendi, Section 10.3).*

The value c is in principle chosen to be small (converging to zero as n increases), and corresponds to selecting a value i for the rank beyond which the function f levels out. This is precisely an "elbow selection rule" as is usual with scree-plots in PCA. The analogy with traditional scree plots and PCA is, in fact, quite strong: in traditional PCA, for each i one determines a rank i matrix that best fits the empirical covariance, and then chooses an appropriate i via a scree plot. Here we do *almost that*: for each i, we determine a rank i matrix that best fits the band-deleted empirical covariance, and then we choose an appropriate i via a scree plot. Particularly in our case, a clear motivation for the "elbow" approach comes from the fact that if we could solve 7.2 with R^K instead of R_n^K, then we would have:

$$f(i) > 0 \text{ if } i = 1, \ldots, r-1, \quad \text{and} \quad f(i) = 0 \text{ if } i \geq r.$$

The asymptotic validity of this motivation is shown in the Section 10.3 of the Appendix.

Going back to Step (A), another difference with traditional PCA, is that the best rank i approximation of the off-band elements of the empirical covariance cannot be determined in closed form by simple eigenanalysis. Thus, we must use approximate schemes in order to solve the optimisation problem 7.1. For a given value of i, we use the fact that any $K \times K$ positive semi-definite matrix of rank at most i can be factorised as $CC^\top$, with $C \in \mathbb{R}^{K \times i}$. The problem thus reduces to

$$\min_{C \in \mathbb{R}^{K \times i}} \quad \left\| P^K \circ (R_n^K - CC^\top) \right\|_F^2, \tag{7.2}$$

for $i = 1, \ldots, K/4 - 1$. Notice that these problems are *not* convex in C, and we thus do not have guarantees that gradient descent-type algorithms will converge to a global optimum (of which there are multiple, since the matrix factorisation is not unique). That being said, recent theoretical progress (e.g. Chen and Wainwright [4]) shows that, remarkably, projected gradient descent methods with a reasonable starting point have high probability of yielding "good" local optima in factorised matrix completion problems. In our own

implementations (e.g. in our Simulations in Section 8) we solve the optimisation problem 7.2 (which can be seen as factorised matrix completion) using the function `fminunc` of the optimization toolbox in MATLAB [14], with starting point $C_0 = U_i \Sigma_i^{1/2}$, where: $U\Sigma U^T$ is the singular value decomposition of R_n^K; U_i is the $n \times i$ matrix obtained by keeping the first i columns of U; and Σ_i is the $i \times i$ matrix obtained by keeping the first i lines and columns of Σ. This function uses a subspace trust-region method based on the interior-reflective Newton method described in [6] and [5] to perform the optimization. Though we do not use the exact same method, we are in a similar setup as Chen and Wainwright [4], so we can expect to obtain "good" local optima. Indeed, in our Simulations (Section 8) the computational method was stable and quickly converged to a reasonable local optimum.

With $\hat{L}_n^K$ at hand, the estimator $\hat{B}_n^K$ can be calculated as the alternated projection of $\Delta_n^K = R_n^K - \hat{L}_n^K$ onto the intersection of the convex sets of $K \times K$ banded matrices with bandwidth at most $\lceil K/4 \rceil$, and of non-negative $K \times K$ matrices. While there is no closed form for this projection, we can iteratively approximate it either using iterated projections onto each of these sets (directly following the formal definition), or using Dykstra's algorithm (Boyle & Dykstra [3]).

Sample R and Matlab Code for the implementation of our methodology can be found at `http://smat.epfl.ch/code/FDA_MatrixCompletion.zip`.

8. Simulation Study

In order to study the performance of our method on a broad range of setups, we consider nine general scenarios to simulate our data. For each of these scenarios, we simulate n iid mean-zero functions Y_i and n iid mean-zero functions W_i on a grid of K equally spaced points on the interval $[0, 1]$. From these samples of discretised curves we calculate the matrices L_n^K and B_n^K:

$$L_n^K(a,b) = \frac{1}{n}\sum_{i=1}^{n} Y_i(t_a)Y_i(t_b) \quad \text{and} \quad B_n^K(a,b) = \frac{1}{n}\sum_{i=1}^{n} W_i(t_a)W_i(t_b),$$

for $a, b \in \{1, \ldots, K\}$, and then set $R_n^K = L_n^K + B_n^K$.

We construct the smooth curves Y_i by setting $Y_i(t_j) = \sum_{a=1}^{r} c_{ia}\lambda_a^{1/2}\eta_a(t_j)$, where $\lambda_1, \ldots, \lambda_r$ are positive scalars and $c_{ia} \sim N(0,1)$. We consider three different cases for the functions $\eta_1, \ldots, \eta_r$ (which are, by construction, the eigenfunctions of $\mathscr{L}$). In the first case, we take the $\{\eta_j\}_{j=1}^r$ as the first r Fourier basis elements (denoted by FB in the sequel), and for the particular case $r = 1$, instead of using the constant function $\eta_1(t) = 1$, we take $\eta_1(t) = \sin(2\pi t)$; in the second case, the $\{\eta_j\}_{j=1}^r$ are constructed as the Gram-Schmidt orthogonalisation of the first r analytic functions (denoted by AC in the sequel) from the following list:

$$\begin{aligned} &\eta_1(t) = 5t\sin(2\pi t), && \eta_2(t) = t\cos(2\pi t) - 3, && \eta_3(t) = 5t + \sin(2\pi t) - 2, \\ &\eta_4(t) = \cos(4\pi t) + (t/2)^2, && \eta_5(t) = \tfrac{\Gamma(4)}{\Gamma(2)\Gamma(2)} t(1-t). \end{aligned}$$

Finally, in the third case, we take the $\{\eta_j\}_{j=1}^r$ as the first r shifted Legendre polynomials $\tilde{P}_i(x)$ (denoted by LP in the sequel) defined as :

$$\begin{aligned} &\eta_1(t) = 6t^2 - 6t + 1, && \eta_2(t) = 2t - 1, && \eta_3(t) = 1, \\ &\eta_4(t) = 20t^3 - 30t^2 + 12t - 1, && \eta_5(t) = 70t^4 - 140t^3 + 90t^2 - 20t + 1. \end{aligned}$$

The rough curves W_i are produced in one of the following three ways:

1. We set $W_i(t_j) = \sum_{a=0}^{q} \theta_a \varepsilon_{i,j-a}$, where $q = \lceil K\delta/2 \rceil$, $\theta_0 = 1$, $\theta_1 \ldots, \theta_q \in (-1, 1)$ are scalars and $\varepsilon_{i,j} \overset{iid}{\sim} N(0,1)$ (denoted by MA in the sequel).
2. We set $W_i(t_j) = \sum_{a=1}^{d} b_{ia}\beta_a^{1/2}\psi_a(t_j)$, where $\beta_1, \ldots, \beta_d$ are positive scalars and $b_{ia} \sim N(0,1)$. The functions ψ_a are triangular functions of norm 1 with support $[(a-1)\delta, a\delta]$ (denoted by TRI in the sequel).
3. We set $W_i(t_j) = \sum_{a=1}^{d} b_{ia}\beta_a^{1/2}\psi_a(t_j)$, where $\beta_1, \ldots, \beta_d$ are positive scalars and $b_{ia} \sim N(0,1)$. The functions ψ_a are realisations of reflected Brownian bridges defined on $[(a-1)\delta, a\delta]$ (denoted by RBB in the sequel).

The nine different scenarios resulting from the three possible choices for the eigenfunctions η and the three possible choices for the rough component W are summarised in Table 1.

Scenarios	A	B	C	D	E	F	G	H	I
Y_i	FB	AC	LP	FB	AC	LP	FB	AC	LP
W_i	MA	MA	MA	TRI	TRI	TRI	RBB	RBB	RBB

TABLE 1
Scenarios for the simulation study.

For each scenario, we consider 6 different combinations of the rank and bandwidth parameters r and δ, as given in the Table 2.

Combination	1	2	3	4	5	6
r	1	1	3	3	5	5
δ	0.05	0.1	0.05	0.1	0.05	0.1

TABLE 2
Different values of the rank and bandwidth parameter.

Finally, we also consider two different regimes for the choice of the eigenvalues $\lambda_1 < \ldots < \lambda_r$ of $\mathscr{L}$ and $\beta_1 < \ldots < \beta_d$ of $\mathscr{B}$; the first one can be seen as the easy case where there is a clear ordering distinction between the two sets, i.e. $\lambda_r \gg \beta_1$ (regime 1); the second one is the interlaced case, when $\lambda_r < \beta_1 < \lambda_{r-1}$ (regime 2). In regime 1, the r eigenvalues λ are equally spaced between $\lambda_1 = 1.45$ and $\lambda_r = 0.25$, and we use $\lambda_1 = 0.25$ for $r = 1$. In regime 2, the eigenvalues $\{\lambda_1, \ldots, \lambda_r\}$ are equally spaced between $\lambda_1 = 1$ and $\lambda_r = 0.04$. In both regimes, the rough processes are simulated with $\beta_1 = 0.09$ and with the remaining eigenvalues being smaller than 0.04 and decreasing toward zero. For each combination (r, δ) with $r > 1$ of Table 2, we consider each of the two regimes and for the particular case $r = 1$, we consider only regime 1. In total, we consider 10 different cases in each one of the nine simulation scenarios.

Our simulation study is divided into two parts. We first illustrate how the scree plots used to select the rank r of the operator $\mathscr{L}$ behave for the different scenarios. These show that using the scree plot as a basis for selection can be a very reasonable approach. We then compare our estimator $\hat{L}_n^K$ of L_n^K to the one obtained by three other methods: a direct use of a truncated Karhunen-Loève expansion; the spline smoothing approach popularised by Ramsay and Silverman [18]; and the PACE method of Yao et al. [22]. We also construct the estimated predictors $\hat{Y}_n^K$ of Y^K for a subset of the scenarios in order to probe their predictive accuracy. In doing this, we use the true rank of $\mathscr{L}$, as the simulations are computationally very intensive, and it would be infeasible to use an automatic selection method (and of course, it would be impossible to make a choice based on inspection of scree plots for all replications). Note that for the rest of this section we consider the maximal bandwidth of B^K to be 10 instead of $K/4 = 25$ (without emphasising it by a new notation), since one would rarely expect a rough process to have such a long memory, and since using a smaller maximal bandwidth value gives more stable and accurate numerical results.

8.1. Rank Selection

In order to probe the appropriateness of using a scree-type plot in order to estimate the rank r of the operator $\mathscr{L}$, we ran simulations on one sample of each scenario, each combination of the parameters r and δ and both regimes (for a total of $9 \times 6 + 9 \times 4 = 90$ simulations). As explained in Section 7, we plot the function $f(i) = \|P^K \circ (R_n^K - \hat{C}_i \hat{C}_i^\top)\|_F^2$, where $\hat{C}_i \in \mathbb{R}^{K \times i}$ is the minimiser of the optimisation problem 7.2, and then we select the rank j beyond which $f(j)$ levels out, i.e. beyond which no meaningful reduction to the objective function is achieved. In practice we evaluate the function f over $i = 1, \ldots, 10$ and not over $1, \ldots, K/4 - 1 = 24$ as mentioned in the theory since the procedure is quite computationally intensive; it is clear from the resulting plots that this is not restrictive. The results are presented by scenario and by regime in Figure 2. Since the functions f are not on the same scale for every regime and every combination, we plotted a normalised version of f given by $f(i)/\|P^K \circ R_n^K\|$. For each scenario, the function f for the samples generated with $r = 5$ are in black, the ones generated with $r = 3$ are in red and the ones generated with $r = 1$ in blue. The dotted vertical lines indicate the location of the true rank, i.e. 5 (in black), 3 (in red)

and 1 (in blue). The figure reveals that for most of the scenarios, we would select the rank quite accurately in regime 1 and we would underestimate it a little bit in regime 2. In further simulations (reported in the Appendix, Section 10.5) we study the effect of rank misspecification. It seems that underestimation is quite impactful in Regime 1 (non-interlaced eigenvalues) and that overestimation does not have a severe impact in both regimes, which suggests that one should not hesitate to over-estimate the rank relative to what the scree-plot indicates.

8.2. Comparisons

We investigate the performance of our estimator of L_n^K, alongside the three following methods:

1. The spline smoothing approach, popularised by Ramsay and Silverman [18]: compute $\widetilde{X}_i$, the smooth version of the observed curves X_i, by using B-spline smoothing; then define the estimator of L_n^K as $\hat{L}_{RS}^K(a,b) = \frac{1}{n}\sum_{i=1}^n \widetilde{X}_i(t_a)\widetilde{X}_i(t_b)$;
2. The PACE method (Yao et al. [22]) described in Section 1: the estimator of L_n^K is given by $\hat{L}_{PACE}^K(a,b) = \widetilde{\rho}(t_a,t_b)$. Of course it must be noted that PACE was primarily introduced for the sparse sampling case, but it can still be used in a dense setting.
3. Truncation of the empirical Karhunen-Loève (KL) expansion: we derive the spectral decomposition of R_n^K, and the estimator of L_n^K is simply equal to a spectrally truncated version thereof, at a level rk, where rk is chosen such that the variance explained is at least 95%.

For every choice of scenario (A–I), rank/bandwidth combination (1–6), and eigenvalue regime (regime 1 or regime 2), we simulate 100 replications for a sample size of $n = 300$ on a grid of $K = 100$ points. Results for different values of n and K can be found in the Appendix, Section 10.5. For each replicate, we determine the estimators given by the four different methods, and calculate their normalised error, by evaluating the function $\mathrm{Err}(u) = (\|u - L_n^K\|_F)/\|L_n^K\|_F$ at every one of these estimators. We then form the ratio between our method's relative error (in the denominator) and the relative error of each of the three other methods (in the numerator). Consequently, we calculate 3×100 ratios per simulation regime. Their corresponding first quartiles, medians and third quartiles are presented in Table 4, where those medians exceeding 1 have been highlighted in bold. These indicate settings where our approach typically performs comparably or at least as well any as the approach it is being compared to.

Of course, one cannot expect there to be a uniformly best method (for instance, the KL expansion is expected to perform best when all the eigenfunctions are approximately mutually orthogonal and the eigenvalues are not interlaced). That being said, Table 4 reveals that our method has a performance that is typically better than or comparable to that of the best competitor in all but one scenarios/combinations. The exceptional case corresponds to a situation where the smooth curves were generated with the first 5 Legendre polynomials. In this particular setup, our optimisation problem was quite unstable due to the particular shape of the matrix L_n^K – it had very high values on the band relative to values outside the band, rendering matrix completion difficult. Consequently, some of the replications returned estimators that where completely off, as is indicated in the table by the small values of the first quartile for the scenarios C, F and I with $r = 5$. Of course, all the results need to be taken with a grain of salt, as we make use of the true rank when constructing our estimator, which in practice is unknown and must be selected (and of course, the methods to which we compare also involve the choice of tuning parameters, depending on which their performance may vary).

In practice, it may of course be that the rough component is indeed pure noise. In order to check whether our method performs comparably well with the other methods in this more classical setup, we additionally consider a scenario where the smooth curves are generated using a Fourier basis and the rough curves are discrete white noise. In this situation, the matrix B^K representing the discretised kernel b is precisely diagonal instead of just banded. The results are presented in the Table 5. Surprisingly, it appears that our method performs equally well or better than all other methods in all scenarios considered. A likely explanation is that, even when the process W has a diagonal kernel, its finite sample empirical kernel will not be exactly diagonal, but banded (since some empirical correlations will exist).

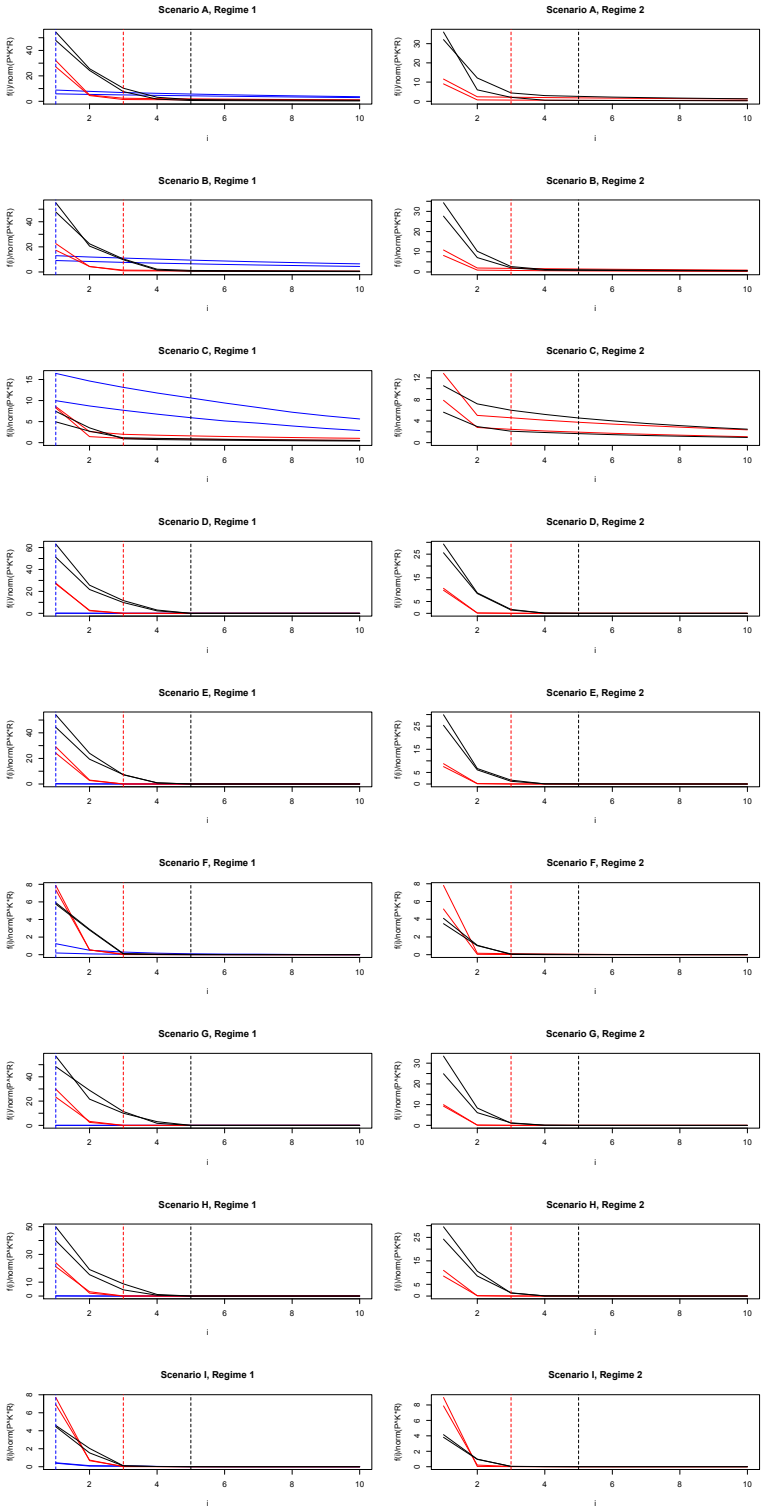
Scenario A, Regime 1
Scenario A, Regime 2
Scenario B, Regime 1
Scenario B, Regime 2
Scenario C, Regime 1
Scenario C, Regime 2
Scenario D, Regime 1
Scenario D, Regime 2
Scenario E, Regime 1
Scenario E, Regime 2
Scenario F, Regime 1
Scenario F, Regime 2
Scenario G, Regime 1
Scenario G, Regime 2
Scenario H, Regime 1
Scenario H, Regime 2
Scenario I, Regime 1
Scenario I, Regime 2

FIG 2. *Plots of the function $f(\cdot)$ (defined in Section 7) normalised by $\|P^K \circ R_n^K\|$ for a given scenario, a given combination of parameters and a given regime. The curves in black correspond to a setting with $r = 5$, those in red to a setting with $r = 3$ and those in blue to a setting with $r = 1$.*

Regime 1				
Scenario	(rk,δ)	PACE	KL	RS
A	$(1, 0.05)$	**4.01** $(2.51, 6.46)$	**2.87** $(1.93, 4.18)$	**4.15** $(3.59, 5.16)$
	$(1, 0.10)$	**4.44** $(2.21, 7.83)$	**3.40** $(2.01, 5.48)$	**4.92** $(3.79, 6.03)$
	$(3, 0.05)$	**3.19** $(2.31, 4.60)$	**2.89** $(2.19, 3.73)$	**3.02** $(2.59, 3.46)$
	$(3, 0.10)$	**3.10** $(2.13, 4.50)$	**2.75** $(1.89, 3.97)$	**2.89** $(2.40, 3.32)$
	$(5, 0.05)$	**2.58** $(2.07, 3.26)$	**2.41** $(2.04, 2.92)$	**2.04** $(1.81, 2.33)$
	$(5, 0.10)$	**2.20** $(1.79, 2.86)$	**2.10** $(1.71, 2.60)$	**1.87** $(1.60, 2.08)$
B	$(1, 0.05)$	**3.95** $(2.05, 5.80)$	**3.09** $(1.79, 4.46)$	**4.30** $(3.51, 5.02)$
	$(1, 0.10)$	**3.54** $(1.83, 6.12)$	**2.55** $(1.55, 4.83)$	**4.18** $(3.44, 5.08)$
	$(3, 0.05)$	**2.93** $(2.55, 4.11)$	**2.85** $(2.36, 3.55)$	**2.72** $(2.37, 3.03)$
	$(3, 0.10)$	**3.16** $(2.49, 4.14)$	**2.74** $(2.22, 3.46)$	**2.71** $(2.43, 3.11)$
	$(5, 0.05)$	**1.91** $(1.49, 2.83)$	**1.84** $(1.48, 2.38)$	**1.49** $(1.23, 1.72)$
	$(5, 0.10)$	**1.62** $(1.28, 2.20)$	**1.57** $(1.25, 2.03)$	**1.35** $(1.07, 1.61)$
C	$(1, 0.05)$	**2.22** $(0.87, 4.20)$	**1.05** $(0.49, 2.27)$	**2.82** $(2.17, 3.71)$
	$(1, 0.10)$	**1.34** $(0.71, 3.02)$	0.63 $(3.38, 1.95)$	**2.23** $(1.01, 3.78)$
	$(3, 0.05)$	**2.08** $(1.58, 2.90)$	**1.73** $(1.28, 2.27)$	**2.19** $(1.78, 2.59)$
	$(3, 0.10)$	**1.52** $(1.08, 2.36)$	**1.33** $(0.79, 2.01)$	**1.95** $(1.33, 2.45)$
	$(5, 0.05)$	0.43 $(0.37, 0.55)$	0.5 $(0.48, 0.75)$	0.42 $(0.28, 0.74)$
	$(5, 0.10)$	0.49 $(0.40, 0.70)$	0.51 $(0.48, 0.69)$	0.44 $(0.28, 0.74)$
D	$(1, 0.05)$	**11.7** $(9.89, 12.8)$	**11.7** $(9.89, 12.8)$	**10.5** $(8.77, 11.6)$
	$(1, 0.10)$	**21.0** $(18.3, 26.5)$	**21.9** $(18.2, 26.4)$	**16.1** $(13.4, 19.3)$
	$(3, 0.05)$	**6.83** $(5.98, 7.41)$	**6.66** $(5.85, 7.33)$	**5.00** $(5.21, 6.46)$
	$(3, 0.10)$	**11.2** $(9.62, 12.9)$	**10.8** $(9.10, 12.4)$	**8.80** $(7.34, 10.0)$
	$(5, 0.05)$	**4.51** $(3.91, 5.18)$	**4.27** $(3.68, 4.95)$	**3.92** $(3.38, 4.52)$
	$(5, 0.10)$	**7.50** $(6.20, 8.65)$	**7.11** $(5.65, 8.24)$	**5.94** $(4.88, 6.74)$
E	$(1, 0.05)$	**7.77** $(6.97, 9.13)$	**7.76** $(6.97, 9.12)$	**7.03** $(6.17, 8.01)$
	$(1, 0.10)$	**15.1** $(12.6, 18.0)$	**15.0** $(12.6, 18.0)$	**11.0** $(9.41, 13.4)$
	$(3, 0.05)$	**5.55** $(5.05, 6.31)$	**5.73** $(5.15, 6.61)$	**4.88** $(4.45, 5.60)$
	$(3, 0.10)$	**9.15** $(7.81, 10.7)$	**9.36** $(8.00, 11.0)$	**7.08** $(5.98, 8.25)$
	$(5, 0.05)$	**2.83** $(2.26, 3.62)$	**3.03** $(2.39, 3.95)$	**2.54** $(1.95, 3.12)$
	$(5, 0.10)$	**5.40** $(4.31, 6.71)$	**5.55** $(4.56, 7.09)$	**4.30** $(3.34, 5.30)$
F	$(1, 0.05)$	**8.91** $(7.56, 10.2)$	**9.05** $(7.69, 10.3)$	**7.78** $(6.77, 9.08)$
	$(1, 0.10)$	**18.2** $(14.6, 24.5)$	**18.3** $(14.7, 24.6)$	**13.3** $(10.9, 17.9)$
	$(3, 0.05)$	**5.43** $(4.58, 6.31)$	**5.67** $(4.82, 6.67)$	**4.69** $(3.89, 5.51)$
	$(3, 0.10)$	**9.84** $(8.83, 11.2)$	**10.2** $(9.12, 11.5)$	**7.47** $(6.51, 8.43)$
	$(5, 0.05)$	0.51 $(0.18, 0.86)$	0.52 $(0.19, 0.91)$	0.44 $(0.15, 0.72)$
	$(5, 0.10)$	**1.03** $(0.47, 2.11)$	**1.07** $(0.49, 2.20)$	0.73 $(0.36, 1.57)$
G	$(1, 0.05)$	**13.5** $(10.2, 17.0)$	**13.4** $(10.2, 16.8)$	**12.1** $(9.43, 15.0)$
	$(1, 0.10)$	**17.2** $(13.0, 24.6)$	**17.2** $(13.0, 25.0)$	**15.6** $(11.6, 20.9)$
	$(3, 0.05)$	**9.78** $(8.17, 11.8)$	**9.21** $(7.38, 11.2)$	**7.93** $(6.97, 9.71)$
	$(3, 0.10)$	**9.76** $(7.94, 12.2)$	**9.34** $(7.58, 12.2)$	**8.64** $(7.00, 10.7)$
	$(5, 0.05)$	**7.05** $(6.07, 8.36)$	**7.15** $(5.93, 8.67)$	**5.64** $(4.85, 7.23)$
	$(5, 0.10)$	**6.93** $(5.68, 8.23)$	**6.44** $(5.37, 8.03)$	**6.00** $(5.04, 7.46)$
H	$(1, 0.05)$	**11.0** $(8.29, 13.8)$	**10.9** $(8.49, 13.8)$	**9.29** $(7.66, 11.8)$
	$(1, 0.10)$	**14.2** $(10.4, 18.2)$	**14.2** $(10.5, 18.2)$	**11.7** $(8.96, 16.0)$
	$(3, 0.05)$	**7.76** $(6.74, 9.89)$	**8.72** $(7.00, 10.2)$	**6.89** $(5.65, 7.85)$
	$(3, 0.10)$	**8.67** $(6.83, 11.2)$	**8.63** $(6.88, 11.3)$	**7.95** $(6.19, 10.2)$
	$(5, 0.05)$	**4.80** $(3.41, 6.20)$	**6.01** $(4.49, 8.14)$	**4.03** $(2.94, 5.44)$
	$(5, 0.10)$	**5.36** $(3.82, 6.89)$	**5.60** $(3.89, 7.17)$	**4.67** $(3.38, 5.95)$
I	$(1, 0.05)$	**11.1** $(9.31, 13.7)$	**11.7** $(9.68, 14.2)$	**9.87** $(8.21, 12.4)$
	$(1, 0.10)$	**16.0** $(11.4, 20.6)$	**16.2** $(11.5, 20.7)$	**13.8** $(9.87, 17.4)$
	$(3, 0.05)$	**7.13** $(6.00, 9.25)$	**7.61** $(6.49, 10.0)$	**6.03** $(5.21, 7.29)$
	$(3, 0.10)$	**7.72** $(6.29, 9.58)$	**8.17** $(6.49, 9.99)$	**6.76** $(5.46, 8.43)$
	$(5, 0.05)$	**1.06** $(0.65, 1.53)$	**1.33** $(0.72, 1.92)$	0.88 $(0.53, 1.27)$
	$(5, 0.10)$	0.94 $(0.18, 1.77)$	0.99 $(0.19, 1.82)$	0.78 $(0.15, 1.54)$

TABLE 3

Table containing the median (the first and third quartiles are in parentheses) of the ratios for the three methods we compared our method with and for the 9 scenarios we considered with the regime 1. We highlight in bold the medians that exceed 1.

Regime 2				
Scenario	Combination	PACE	KL	RS
A	$(3, 0.05)$	**1.84** $(1.16, 2.54)$	**1.87** $(1.09, 2.79)$	**2.26** $(1.28, 2.86)$
	$(3, 0.10)$	**1.20** $(0.95, 1.87)$	0.98 $(0.83, 1.89)$	**1.14** $(0.78, 2.17)$
	$(5, 0.05)$	**1.06** $(0.87, 1.61)$	0.96 $(0.86, 1.72)$	**1.08** $(0.62, 1.76)$
	$(5, 0.10)$	**1.01** $(0.84, 1.24)$	0.93 $(0.82, 1.25)$	0.91 $(0.63, 1.22)$
B	$(3, 0.05)$	**2.11** $(1.29, 2.90)$	**2.22** $(1.22, 2.95)$	**2.06** $(1.30, 2.65)$
	$(3, 0.10)$	**1.32** $(1.05, 1.78)$	**1.10** $(0.91, 1.73)$	**1.26** $(0.73, 2.24)$
	$(5, 0.05)$	0.94 $(0.82, 1.10)$	0.89 $(0.80, 1.04)$	0.75 $(0.44, 1.07)$
	$(5, 0.10)$	**1.04** $(0.87, 1.24)$	0.94 $(0.80, 1.17)$	0.90 $(0.60, 1.33)$
C	$(3, 0.05)$	**1.18** $(0.88, 1.61)$	0.80 $(0.64, 1.44)$	**1.25** $(0.93, 2.02)$
	$(3, 0.10)$	**1.15** $(0.85, 1.62)$	0.72 $(0.58, 1.53)$	**1.35** $(0.83, 1.91)$
	$(5, 0.05)$	0.68 $(0.54, 0.89)$	0.53 $(0.48, 0.71)$	0.79 $(0.52, 1.32)$
	$(5, 0.10)$	0.74 $(0.54, 1.03)$	0.56 $(0.47, 1.04)$	0.77 $(0.58, 1.26)$
D	$(3, 0.05)$	**5.70** $(5.06, 6.62)$	**5.59** $(5.03, 6.65)$	**4.93** $(4.42, 5.73)$
	$(3, 0.10)$	**10.7** $(8.66, 12.2)$	**10.5** $(8.48, 12.2)$	**8.03** $(6.39, 9.37)$
	$(5, 0.05)$	**3.58** $(3.10, 4.18)$	**3.48** $(3.05, 4.03)$	**3.08** $(2.73, 3.59)$
	$(5, 0.10)$	**6.81** $(5.64, 8.09)$	**6.63** $(5.54, 7.72)$	**5.27** $(4.23, 6.17)$
E	$(3, 0.05)$	**4.60** $(3.89, 5.43)$	**4.66** $(3.96, 5.45)$	**4.16** $(3.60, 4.81)$
	$(3, 0.10)$	**8.59** $(6.96, 10.2)$	**8.65** $(7.00, 10.2)$	**6.51** $(5.22, 7.80)$
	$(5, 0.05)$	**2.09** $(1.11, 2.76)$	**2.14** $(1.13, 2.82)$	**1.84** $(0.94, 2.45)$
	$(5, 0.10)$	**3.96** $(3.15, 5.46)$	**4.24** $(3.33, 5.72)$	**3.12** $(2.42, 4.27)$
F	$(3, 0.05)$	**1.13** $(0.06, 2.74)$	**1.17** $(0.07, 2.83)$	0.99 $(0.06, 2.47)$
	$(3, 0.10)$	**3.45** $(0.16, 7.03)$	**3.55** $(0.16, 7.20)$	**2.61** $(0.11, 5.21)$
	$(5, 0.05)$	0.78 $(0.07, 1.43)$	0.81 $(0.07, 1.50)$	0.66 $(0.06, 1.27)$
	$(5, 0.10)$	0.70 $(0.09, 2.85)$	0.71 $(0.09, 2.95)$	0.52 $(0.07, 2.13)$
G	$(3, 0.05)$	**7.87** $(6.60, 9.69)$	**7.31** $(6.22, 9.55)$	**6.56** $(5.56, 8.07)$
	$(3, 0.10)$	**8.05** $(6.46, 9.91)$	**8.02** $(6.41, 9.92)$	**7.03** $(5.58, 9.10)$
	$(5, 0.05)$	**5.73** $(4.73, 6.52)$	**7.03** $(5.95, 8.53)$	**4.94** $(3.92, 5.68)$
	$(5, 0.10)$	**5.87** $(4.77, 7.88)$	**5.75** $(4.69, 7.92)$	**5.30** $(4.35, 7.00)$
H	$(3, 0.05)$	**7.10** $(6.07, 8.22)$	**6.99** $(5.73, 8.16)$	**6.06** $(5.13, 7.17)$
	$(3, 0.10)$	**7.51** $(6.03, 9.43)$	**7.61** $(6.09, 9.53)$	**6.74** $(5.63, 8.19)$
	$(5, 0.05)$	**3.84** $(3.16, 4.91)$	**5.26** $(4.11, 6.90)$	**3.40** $(2.64, 4.14)$
	$(5, 0.10)$	**3.89** $(1.76, 5.46)$	**4.30** $(1.82, 5.84)$	**3.53** $(1.47, 5.02)$
I	$(3, 0.05)$	**4.94** $(3.27, 6.13)$	**5.32** $(3.48, 6.54)$	**4.41** $(3.12, 5.30)$
	$(3, 0.10)$	**3.11** $(0.20, 6.11)$	**3.16** $(0.20, 6.24)$	**2.87** $(0.17, 5.12)$
	$(5, 0.05)$	0.59 $(0.06, 1.47)$	0.67 $(0.07, 1.58)$	0.49 $(0.05, 1.24)$
	$(5, 0.10)$	**1.16** $(0.14, 2.54)$	**1.20** $(0.15, 2.60)$	**1.02** $(0.11, 2.38)$

TABLE 4

Table containing the median (the first and third quartiles are in parentheses) of the ratios for the three methods we compared our method with and for the 9 scenarios we considered with the regime 2. We highlight in bold the medians that exceed 1*.*

Regime 1			
r	PACE	KL	RS
1	**1.92** $(1.69, 2.16)$	**1.76** $(1.53, 2.05)$	**4.08** $(3.77, 4.33)$
3	**2.90** $(2.58, 3.16)$	**3.02** $(2.66, 3.28)$	**3.36** $(3.05, 3.53)$
5	**2.80** $(2.61, 3.01)$	**2.78** $(2.56, 3.02)$	**2.40** $(2.24, 2.65)$
Regime 2			
r	PACE	KL	RS
3	**1.63** $(1.45, 1.76)$	**2.01** $(1.85, 2.19)$	**2.36** $(2.22, 2.57)$
5	**1.28** $(1.16, 1.37)$	**1.48** $(1.36, 1.61)$	**1.64** $(1.45, 1.75)$

TABLE 5

Table containing the median (the first and third quartiles are in parentheses) of the ratios for the three methods we compared our method with for the classical scenario where the rough component is a white noise. We highlight in bold the results that exceed 1*.*

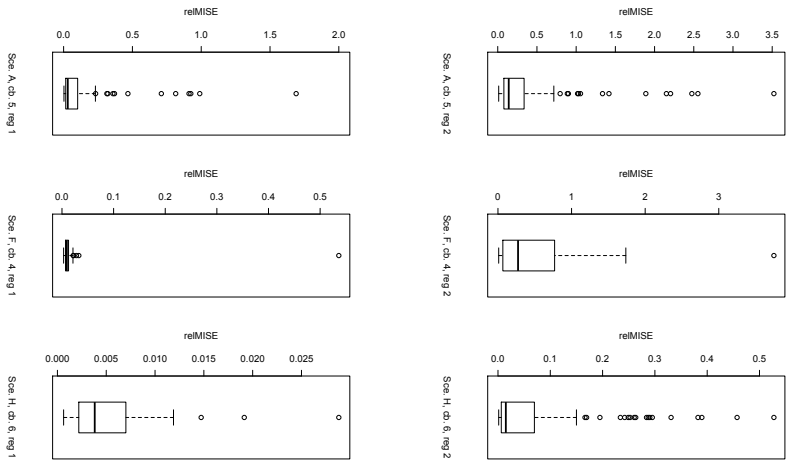

FIG 3. *Distributions of* relMISE. *First row : scenario A with* $r = 5, \delta = 0.05$*, regime 1 on the left and regime 2 on the right. Middle row : scenario F with* $r = 3, \delta = 0.1$*, regime 1 on the left and regime 2 on the right. Last row : scenario H with* $r = 5, \delta = 0.1$*, regime 1 on the left and regime 2 on the right.*

8.3. Prediction of the smooth curves

We selected 6 different cases in order to probe the performance of our estimated predictor $\hat{Y}_n^K$ as a proxy for the true predictor $\Pi(X^K)$. We considered, for both regimes, combination 5 of scenario A, combination 4 of scenarios F and combination 6 of scenario H. For every sample, we calculated the average of the approximation of the normalised mean integrated squared error of $\hat{Y}_n^K$:

$$\text{relMISE} = \frac{1}{n}\sum_{i=1}^{n}\frac{\sum_{j=1}^{K}[\hat{Y}_{n,i}^K(t_j) - \Pi(X_i^K)(t_j)]^2}{\sum_{j=1}^{K}[\Pi(X_i^K)(t_j)]^2}.$$

Figure 3 contains boxplots of their distributions. These illustrate that, as expected, our predictions perform better when the eigenvalues of $\mathscr{L}$ and $\mathscr{B}$ are not interlaced.

9. Proofs of Formal Statements

Proofs of Theorems in Section 3

Proof of Theorem 1. Since the eigenfunctions of $\mathscr{L}_1$ and $\mathscr{L}_2$ are analytic and $\max\{r_1, r_2\} < \infty$, it follows that the corresponding covariance kernels are bivariate analytic functions on $[0,1]^2$ (Krantz and Parks [12, Thm 4.3.3]).

This being the case, the zero set of either kernel is at most 1-dimensional, unless the kernels are uniformly zero (Krantz and Parks [12, Thm 6.33]). Since our theorem follows trivially if $\mathscr{L}_1$ and $\mathscr{L}_2$ are the zero operator, we can assume that their kernels are not uniformly zero. Thus, if we can show that the two kernels coincide on an open subset U of $[0,1]^2$, then they will necessarily coincide everywhere on $(0,1)^2$, and thus on $[0,1]^2$ by continuity. This, in particular, will in turn imply that $\mathscr{B}_1$ and $\mathscr{B}_2$ also coincide.

Without lost of generality, assume that $\delta_1 \geq \delta_2$. Define

$$U = \left(1 - \frac{(1-\delta_1)}{2}, 1\right) \times \left(1 - \frac{(1-\delta_1)}{2}, 1\right).$$

Since $\mathscr{L}_1 + \mathscr{B}_1 = \mathscr{L}_2 + \mathscr{B}_2$, but $\mathscr{B}_1 = \mathscr{B}_2 = 0$ on U, it must be that the kernels of $\mathscr{L}_1$ and $\mathscr{L}_2$ coincide on the open set U, and the proof is complete. □

Proof of Proposition 1. We will first prove the results referring to the processes Z and Y, and then those referring to their covariances, $\mathscr{G}$ and $\mathscr{L}$. Let μ be the mean function of Z and

$$\mathscr{G} = \sum_{n=1}^{r} \theta_n \phi_n \otimes \phi_n,$$

be the specrtum of $\mathscr{G}$, with $\{\theta_n, \phi_n\}$ the corresponding eigenvalues/eigenfunctions. Now let $\epsilon > 0$ be arbitrary, and define $\gamma = \epsilon/\text{trace}\{\mathscr{G}\}$. Define the function $f_{n,J}$ to be the order J Fourier series approximation of ϕ_n, and note that this is an analytic function for all $J < \infty$ (and of course all n). Since Fourier series are dense in L^2, we know that there exists $J_1, ..., J_r$ such that

$$\|\phi_n - f_{n,J_n}\|_{L^2} < \gamma.$$

In particular, if we we pick $J_* = \max\{J_1, ..., J_r\}$ and define $f_n = f_{n,J^*}$, we have that

$$\sup_{1 \le n \le r} \|\phi_n - f_n\|_{L^2} < \gamma.$$

The functions f_n are, of course, analytic. Finally, define a new random function Y via the random series

$$Y = \mu + \sum_{n=1}^{r} \underbrace{\langle Z - \mu, \phi_n \rangle_{L^2}}_{=\xi_n} f_n = Z + \sum_{n=1}^{r} \xi_n e_n,$$

where $e_n = f_n - \phi_n$ satisfies $\|e_n\|_{L^2} < \gamma$. Note that since the $\{f_n\}_{n=1}^r$ are analytic and finitely many, their span consists of analytic functions. Thus the eigenfunctions of the covariance of Y (which are not necessarily exactly equal to the f_n) are analytic too. Furthermore, the rank of Y can clearly not exceed r, whatever the value of ϵ. Now, since the $\{\xi_n\}$ are mean-zero, uncorrelated, and of variance $\{\theta_n\}$, we may write

$$\begin{aligned} \mathbb{E}\|Z - Y\|_{L^2}^2 &= \mathbb{E}\int_0^1 \left(\sum_{n=1}^{r} \xi_n e(t)\right)^2 dt = \int_0^1 \mathbb{E}\left(\sum_{n=1}^{r} \xi_n e(t)\right)^2 dt = \int_0^1 \sum_{n=1}^{r} \theta_n e_n^2(t) dt \\ &= \sum_{n=1}^{r} \theta_n \|e_n\|_{L^2}^2 < \gamma \,\text{trace}\{\mathscr{G}\} = \epsilon. \end{aligned}$$

If we happen to know that $\{\phi_n\}$ are C^1, we may define again $\gamma = \epsilon/\text{trace}\{\mathscr{G}\}$, but now re-define f_n to be trigonometric functions such that

$$\sup_{1 \le n \le r} \|\phi_n - f_n\|_{\infty} < \gamma^{1/2}.$$

This is possible, since the eigenfunctions $\{\phi_n\}$ are C^1, and thus can be uniformly approximated by Fourier series. Define Y and e_n as before, but with the new definition of f_n in place. Once again, since the $\{\xi_n\}$ are mean-zero, uncorrelated, and of variance $\{\theta_n\}$, we have that for any $t \in [0, 1]$,

$$\mathbb{E}(Z(t) - Y(t))^2 = \mathbb{E}\left[\sum_{n=1}^{r} \xi_n e_n(t)\right]^2 = \sum_{i=1}^{r} \theta_n e_n^2(t) < \gamma \,\text{trace}\{\mathscr{G}\} = \epsilon.$$

Now let us focus on the approximation of $\mathscr{G}$ itself. Let $\epsilon > 0$, and define $\gamma = \epsilon/(2 \cdot \text{trace}\{\mathscr{G}\})$. Write

$$\mathscr{G} = \sum_{n=1}^{r} \theta_n \phi_n \otimes \phi_n,$$

with $\{\theta_n, \phi_n\}$ its eigenvalues/eigenfunctions. Define the function $f_{n,J}$ to be the order J Fourier series approximation of ϕ_n, as before. Again, there exist $J_1, ..., J_r$ such that

$$\|\phi_n - f_{n,J_n}\|_{L^2} < \gamma.$$

Set $J_* = \max\{J_1, ..., J_r\}$ and define $f_n = f_{n,J^*}$, so that

$$\sup_{1\leq n\leq r} \|\phi_n - f_n\|_{L^2} < \gamma.$$

The functions f_n are, of course, analytic. Now define the operator $\mathscr{L}$ to be

$$\mathscr{L} = \sum_{n=1}^{r} \theta_n f_n \otimes f_n.$$

This operator is analytic, and has rank at most r. Furthermore, its eigenfunctions are analytic, since they lie in the span of $\mathscr{L}$, which is spanned by the analytic f_n. We now have:

$$\begin{aligned}
\|\mathscr{G} - \mathscr{L}\|_* &\leq \sum_{n=1}^{r} \theta_n \|\phi_n \otimes \phi_n - f_n \otimes f_n\|_* \\
&= \sum_{n=1}^{r} \theta_n \|\phi_n \otimes \phi_n - \phi_n \otimes f_n + \phi_n \otimes f_n - f_n \otimes f_n\|_* \\
&\leq \sum_{n=1}^{r} \theta_n \left\{\|\phi_n \otimes (\phi_n - f_n)\|_* + \|(\phi_n - f_n) \otimes f_n\|_*\right\} \\
&= \sum_{n=1}^{r} \theta_n \left\{\|\phi_n\|_{L^2}\|\phi_n - f_n\|_{L^2} + \|\phi_n - f_n\|_{L^2}\|f_n\|_{L^2}\right\} \\
&= \sum_{n=1}^{r} \theta_n (1 + \|f_n\|_{L^2})\|\phi_n - f_n\|_{L^2} \\
&< 2\gamma \,\mathrm{trace}\{\mathscr{G}\} = \epsilon
\end{aligned}$$

where we used the fact that $\|f_n\|_{L^2} < 1$. If we know that the eigenfunctions $\{\phi_n\}$ of $\mathscr{G}$ are C^1, the Fourier series expansion of each $\phi_n(t)$ converges *uniformly* and *absolutely*. Let $c_1 < \infty$ be the maximum of the ℓ_1 norms of the Fourier coefficients of $\phi_1, ..., \phi_r$ ($c_1 < \infty$ by absolute convergence of the respective Fourier series). Re-define

$$\gamma = \epsilon \times \left[\left(c_1 + \sup_{1\leq n\leq r} \|\phi_n\|_\infty\right) \mathrm{trace}\{\mathscr{G}\}\right]^{-1}.$$

Following the same steps as before, we can choose a J^* sufficiently large, such that setting $f_n = f_{n,J^*}$ we have

$$\sup_{1\leq n\leq r} \|\phi_n - f_n\|_\infty < \gamma.$$

It now follows that

$$\begin{aligned}
\|g - \ell\|_\infty &\leq \sum_{n=1}^{r} \theta_n \sup_{s,t} |\phi_n(s)\phi_n(t) - f_n(s)f_n(t)| \\
&= \sum_{n=1}^{r} \theta_n \sup_{s,t} |\phi_n(s)\phi_n(t) - \phi_n(s)f_n(t) + \phi_n(s)f_n(t) - f_n(s)f_n(t)| \\
&\leq \sum_{n=1}^{r} \theta_n \left\{\sup_t \sup_s |\phi_n(s)|\,|\phi_n(t) - f_n(t)| + \sup_s \sup_t |f_n(t)|\,|\phi_n(s) - f_n(s)|\right\} \\
&\leq \sum_{n=1}^{r} \theta_n (c_1 + \sup_t |\phi_n(t)|)\|\phi_n - f_n\|_\infty \\
&< \left(c_1 + \sup_{1\leq n\leq r} \|\phi_n\|_\infty\right) \gamma \,\mathrm{trace}\{\mathscr{G}\} = \epsilon.
\end{aligned}$$

Finally, for any $\epsilon > 0$, we can replace the specific truncation $J^*(\epsilon)$ used in each of the four parts of the proof, by the largest of all these $J^*(\epsilon)$, and so ϵ can be chosen to be the same in all the approximation results. This concludes the proof. □

Moving on, the proof of Theorem 2 rests upon the observation that it is essentially a statement regarding matrix completion. Our strategy of proof will thus be to translate our functional conditions on $\mathscr{B}$ and $\mathscr{L}$ into matrix properties of L^K and B^K that suffice for unique matrix completion. We first develop the said matrix properties in the form of Lemma 1 and Theorem 4.

Lemma 1. *Let $b(s,t)$ be a continuous kernel on $[0,1]^2$ such that $b(s,t) = 0$ whenever $|s-t| > \delta$, and let $(t_1, \dots, t_K) \in \mathcal{T}_K$ be a grid of K points. Then, the matrix $B^K = \{b(t_i, t_j)\}_{i,j=1}^K$ is banded with bandwidth $2\lceil \delta \cdot K \rceil + 1$.*

Theorem 4. *Let $\mathscr{L}$ have kernel $\ell(s,t) = \sum_{i=1}^r \lambda_i \eta_i(s)\eta_i(t)$ with $r < \infty$ and real analytic orthonormal eigenfunctions $\{\eta_1, \dots, \eta_r\}$. If $K > r$, then the minors of order r of the matrix $L^K = \{\ell(t_i, t_j)\}_{i,j=1}^K$ are all non-zero, almost everywhere on $\mathcal{T}_K$.*

Proof. First notice that from $\ell(s,t) = \sum_{i=1}^r \lambda_i \eta_i(s)\eta_i(t)$, we have

$$L^K_{jl} = \sum_{i=1}^r \lambda_i \eta_i(t_j)\eta_i(t_l).$$

Thus, L^K can be written as $U^K \Sigma (U^K)^\top$, where

$$U^K = \begin{pmatrix} \eta_1(t_1) & \eta_2(t_1) & \cdots & \eta_r(t_1) \\ \eta_1(t_2) & \eta_2(t_2) & \cdots & \eta_r(t_2) \\ \vdots & \vdots & & \vdots \\ \eta_1(t_K) & \eta_2(t_K) & \cdots & \eta_r(t_K) \end{pmatrix} \quad \text{and} \quad \Sigma = \begin{pmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & \lambda_r \end{pmatrix}. \tag{9.1}$$

Any $r \times r$ submatrix of L^K obtained by deleting rows and columns, can then be written as

$$U^K_F \Sigma (U^K_{F'})^\top,$$

where U^K_F (resp. $U^K_{F'}$) is an $r \times r$ matrix obtained by deleting rows of U^K whose indices are not included in $F \subseteq \{1, \dots, K\}$ (resp. F'). The condition that any minor of order r of L^K be non-zero is then equivalent to the condition that

$$\det\left[U^K_F \Sigma (U^K_{F'})^\top\right] = \det[U^K_F]\det[\Sigma]\det[U^K_{F'}] \neq 0,$$

for any subset $F, F' \subseteq \{1, \dots, K\}$ of cardinality r. By construction $\det(\Sigma) \neq 0$, so the minor condition is then equivalent to requiring that $\det(U^K_F) \neq 0$ for any subset $F \subseteq \{1, \dots, K\}$ of cardinality r.

We will show that this is indeed the case almost everywhere on $\mathcal{T}_K$. Let μ denote Lebesgue measure on $\mathcal{T}_K$ and let $F = \{1, \dots, r\}$, without loss of generality (so that U^K_F is formed by keeping the first r rows of U^K). Using the Leibniz formula, we have that $\det(U^K_F)$ can be written as the function

$$D(t_1, \dots, t_r) = \sum_{\sigma \in S_r} \varepsilon(\sigma) \prod_{i=1}^r \eta_i(t_{\sigma(i)}),$$

where S_r is the symmetric group on r elements and $\varepsilon(\sigma)$ is the signature of the permutation σ. Note that the function D is real analytic on $(0,1)^r$, by virtue of each η_i being real analytic on $(0,1)$.

We will now proceed by contradiction. Assume that

$$\mu\{(x_1, \dots, x_K) \in \mathcal{T}_K : D(x_1, \dots, x_r) = 0\} > 0.$$

Since μ is Lebesgue measure, it follows that the Hausdorff dimension of the set $A = \{(x_1, \dots, x_r) : D(x_1, \dots, x_r) = 0\}$ is equal to r. However, since D is analytic, Krantz and Parks [12, Thm 6.33] implies the

dichotomy: either D is constant everywhere on $(0,1)^r$, or the set A is at most of dimension $r-1$. Thus it must be that D is everywhere constant on $(0,1)^r$, the constant being of course zero:

$$D(x_1,\dots,x_r) = \sum_{\sigma\in S_r} \varepsilon(\sigma) \prod_{i=1}^{r} \eta_i(x_{\sigma(i)}) = 0,\ \forall\ (x_1,\dots,x_r) \in (0,1)^r.$$

Now fix $(x_1,\dots,x_{r-1})$ and apply to D (viewed as a function of x_r only) the continuous linear functional $T_{\eta_r}(f) = \langle f, \eta_r\rangle$. We obtain that for all $(x_1,\dots,x_{r-1}) \in (0,1)^r$:

$$0 = \langle D, \eta_r\rangle \quad = \quad \sum_{\sigma\in S_r} \varepsilon(\sigma) \left[\prod_{i:\sigma(i)\neq r} \eta_i(x_{\sigma(i)})\right] \langle \eta_{\sigma^{-1}(r)}, \eta_r\rangle = \sum_{\sigma\in S_{r-1}} \varepsilon(\sigma) \prod_{i=1}^{r-1} \eta_i(x_{\sigma(i)}).$$

Applying iteratively the continuous linear functionals $T_{\eta_j}(f) = \langle f, \eta_j\rangle$ to D while keeping $(x_1,\dots,x_{j-1})$ fixed then leads to

$$\eta_1(y) = 0,\ \forall\, y \in (0,1).$$

This last equality contradicts the fact that η_1 is of norm one, and allows us to conclude that $\mu\{(x_1,\dots,x_K) \in \mathcal{T}_K : D(x_1,\dots,x_r) = 0\} = 0$. □

We now prove Theorem 2 by demonstrating that the matrix properties of (L^K, B^K) that derive from its assumptions are sufficient for unique matrix completion. The proof is inspired by Proposition 2.12 of [11].

Proof of Theorem 2. Given our conditions, Lemma 1 implies that $B_1, B_2 \in \mathbb{R}^{K\times K}$ are banded matrices with bandwidth $2\lceil \delta_i \cdot K\rceil + 1$, for $i \in \{1,2\}$.

Let $\delta = \max\{\delta_1, \delta_2\}$ and assume without loss of generality that $r_1 \geq r_2$. Let Ω be the set of indices on which both B_1 and B_2 vanish, which by Lemma 1 is $\Omega = \{(i,j) \in \{1,\dots,K\}^2 : |i-j| > \lceil \delta \cdot K\rceil\}$. From $L_1 + B_1 = L_2 + B_2$, we obtain that $\{L_1\}_{ij} = \{L_2\}_{ij}, \forall (i,j) \in \Omega$. Let Ω_A be the set of indices of a submatrix formed by the first r_1 rows and the last r_1 columns of a $K \times K$ matrix, the condition $K \geq K^* = \frac{2r_1+2}{1-2\delta}$ implies that $\Omega_A \subset \Omega$, which in turn implies that the matrices L_1 and L_2 contain a common submatrix A of dimension $r_1 \times r_1$.

Assume that all minors of order r_1 of L_1 are non-zero. Then, the determinant of A is non-zero, which implies that the rank of L_2 is also r_1. We thus establish that L_1 and L_2 are two rank r_1 matrices equal on Ω. Let L^* be a matrix equal to L_1 on Ω, but unknown at those indices that do not belong to Ω. We will now show that there exists a unique rank r_1 completion of L^*. Due to the band pattern of the unobserved entries of L^* and the inequality $K \geq K^* = \frac{2r_1+2}{1-2\delta}$, it is possible to find a submatrix of L^* of dimension $(r_1+1)\times(r_1+1)$ with only one unobserved entry, denoted x^*. Using the fact that the determinant of any square submatrix of dimension bigger than r_1+1 is zero, we obtain a linear equation of the form $ax^* + b = 0$, where a is equal to the determinant of a submatrix of dimension $r_1 \times r_1$. Since we assume that any minor of order r_1 is non-zero, we have that $a \neq 0$ and the previous equation has a unique solution. It is then possible to impute the value of x^*. Applying this procedure iteratively until all missing entries are determined allows us to uniquely complete the matrix L^* into a rank r_1 matrix. In summary, we have demonstrated that when all minors of order r_1 of L_1 are non-zero, it holds that $L^* = L_1 = L_2$ and hence $B_1 = B_2$. Theorem 4 assures us that L_1 indeed has non-vanishing minors of order r_1 almost everywhere on $\mathcal{T}_K$, and so we conclude that it must be that $L_1 = L_2$ and $B_1 = B_2$ almost everywhere on $\mathcal{T}_K$ □

Proofs of Theorems in Section 4

Proof of Proposition 2. Since $\delta < 1/4$ and $K \geq 4r+1$ implies $K \geq \frac{2r+2}{1-2\delta}$, we can use directly Theorem 2 which tells us that L^K is the only matrix of rank at most r such that $P^K \circ R^K = P^K \circ L^K$, and then such that $\|P_K \circ (R^K - L^K)\|_F^2 = 0$. The objective function 4.1 thus achieves its minimal value of r at L^K. We now turn to prove that

$$L^K = \underset{\theta\in\mathbb{R}^{K\times K}}{\arg\min} \left\{ \left\| P^K \circ (R^K - \theta)\right\|_F^2 + \tau\, \mathrm{rank}(\theta)\right\},$$

for all $\tau > 0$ sufficiently small. Since we have established that L^K uniquely solves

$$\min_{\theta \in \mathbb{R}^{K\times K}} \operatorname{rank}\{\theta\} \qquad \text{subject to } \left\|P^K \circ (R^K - \theta)\right\|_F^2 = 0,$$

it follows that for all $\tau > 0$ and any $\theta \in \mathbb{R}^{K\times K}$ of rank greater or equal to r, we have that

$$\left\|P^K \circ (R^K - L^K)\right\|_F^2 + \tau \operatorname{rank}(L^K) < \left\|P^K \circ (R^K - \theta)\right\|_F^2 + \tau \operatorname{rank}(\theta).$$

We thus concentrate on matrices $\theta \in \mathbb{R}^{K\times K}$ of rank at most $r-1$, for $r > 1$. Let

$$\mu = \min_{\theta \in \mathbb{R}^{K\times K},\ \operatorname{rank}(\theta)\le r-1} \{\left\|P^K \circ (R^K - \theta)\right\|_F^2\} > 0.$$

Now let $\tau_* = \frac{\mu}{r-1}$. Then, for any $\tau < \tau_*$, and any θ of rank less than r,

$$\left\|P^K \circ (R^K - L^K)\right\|_F^2 + \tau \operatorname{rank}(L^K) = \tau r < \mu + \tau \le \left\|P^K \circ (R^K - \theta)\right\|_F^2 + \tau \operatorname{rank}(\theta).$$

In summary, putting our results together, we have shown that for all $\tau \in (0, \tau_*)$,

$$L^K = \arg\min_{\theta \in \mathbb{R}^{K\times K}} \left\{ \left\|P^K \circ (R^K - \theta)\right\|_F^2 + \tau \operatorname{rank}(\theta) \right\}.$$

Finally, it is worth pointing out that although τ_* depends on r, this does not mean that the objective function depends on unknowns: r can be shown (using Theorem 4) to be equal to the rank of the submatrix formed by the first $\lceil K/4 \rceil$ rows and the last $\lceil K/4 \rceil$ columns of R^K, and thus we can determine τ_* directly from the matrix R^K. This completes the proof.

□

Proofs of Theorems in Section 6

Proof of Theorem 3. We begin by the usual bias/variance decomposition

$$\begin{aligned} \left\|\hat{\mathscr{L}}_n^K - \mathscr{L}\right\|_{\mathrm{HS}}^2 &\le 2\left\|\hat{\mathscr{L}}_n^K - \mathscr{L}^K\right\|_{\mathrm{HS}}^2 + 2\left\|\mathscr{L}^K - \mathscr{L}\right\|_{\mathrm{HS}}^2 \\ &= 2K^{-2}\left\|\hat{L}_n^K - L^K\right\|_{\mathrm{F}}^2 + 2\left\|\mathscr{L}^K - \mathscr{L}\right\|_{\mathrm{HS}}^2. \end{aligned}$$

For the second term (bias), we note that by a Taylor expansion

$$\begin{aligned} \int_0^1\int_0^1 (\ell(x,y) - \ell_K(x,y))^2 dxdy &= \sum_{i,j=1}^K \int_{I_{i,K}}\int_{I_{j,K}} (\ell(x,y) - \ell(t_i,t_j))^2 dxdy \\ \le \sum_{i,j=1}^K \int_{I_{i,K}}\int_{I_{j,K}} 2K^{-2} \sup_{(x,y)\in I_{i,K}\times I_{j,K}} &\qquad \|\nabla\ell(x,y)\|_2^2 \le 2K^{-2} \sup_{(x,y)\in[0,1]^2} \|\nabla\ell(x,y)\|_2^2. \end{aligned}$$

It now remains to be shown that $K^{-2}\left\|\hat{L}_n^K - L^K\right\|_{\mathrm{F}}^2 = O_{\mathbb{P}}(n^{-1})$, and that the $O_{\mathbb{P}}(n^{-1})$ is uniform in K. To show that this is valid (Lebesgue) almost everywhere on $\mathcal{T}_K$, we will show that it is valid almost surely, when choosing a grid uniformly at random from $\mathcal{T}_K$ (following the discussion in Remark 3). Using the fact that $K = m \times K^*$, $m \ge 1$,we consider a simple coupling construction, as this will be convenient for our proof:

1. Choose $t_{j,K}$ uniformly at random from $I_{j,K}$, independently for all $j = 1, \ldots, K$, writing $\boldsymbol{t}_K = \{t_{j,K}\}_{j=1}^K$.
2. Choose t_{j,K^*} uniformly at random among the m elements of $\boldsymbol{t}_K \cap I_{j,K^*}$, and write $\boldsymbol{t}_{K^*} = \{t_{j,K^*}\}_{j=1}^{K^*}$.
3. Define the set $q = \{j \in \{1,\ldots,K\} : t_{j,K} \in \boldsymbol{t}_{K^*}\}$ of cardinality K^*.

Note that our construction guarantees that $\boldsymbol{t}_{K^*} \subseteq \boldsymbol{t}_K$, with $\boldsymbol{t}_{K^*}$ being marginally uniformly distributed on $\mathcal{T}_{K^*}$. Define Θ_K to be the space of $K \times K$ covariance operators of rank smaller than $K/4$. We consider the functionals

$$\mathbb{S}_{n,K} : \Theta_K \to [0,\infty), \qquad \mathbb{S}_{n,K}(\theta) = \underbrace{K^{-2}\|P^K \circ (\theta - R_n^K)\|_{\mathrm{F}}^2}_{\mathbb{M}_{n,K}(\theta)} + \tau_n K^{-2}\mathrm{rank}(\theta),$$

$$S_{n,K} : \Theta_K \to [0,\infty), \qquad S_{n,K}(\theta) = \underbrace{K^{-2}\|P^K \circ (\theta - R^K)\|_{\mathrm{F}}^2}_{M_K(\theta)} + \tau_n K^{-2}\mathrm{rank}(\theta),$$

where $P_K(i,j) = \mathbf{1}\{|i-j| > \lceil K/4 \rceil\}$. We also define $d_n : \Theta_K \times \Theta_K \to [0,\infty)$ as

$$d_n(\theta_1,\theta_2) = \sqrt{K^{-2}\|\theta_1 - \theta_2\|_F^2 + \tau_n K^{-2}|\mathrm{rank}(\theta_1) - \mathrm{rank}(\theta_2)|},$$

noting that d_n *does not* need to be a metric. We will show that for n sufficiently large, L^K is the unique minimizer of $S_{n,K}(\cdot)$ on Θ_K. Note that, since $K \geq 4r+4$, Theorem 2 implies that L^K is (almost surely) the unique matrix for which the functional $M_K(\theta)$ attains the minimal value of zero on Θ_K^r, where Θ_K^r is defined as the space of $K \times K$ covariance operators of rank at most r. It follows that for any $H \in \Theta_K$ different than L^K and of rank $p \in \{r, \ldots, K/4 - 1\}$,

$$S_{n,K}(L^K) = M_K(L^K) + \frac{r\tau_n}{K^2} = \frac{r\tau_n}{K^2} \leq \frac{p\tau_n}{K^2} < M_K(H) + \frac{p\tau_n}{K^2} = S_{n,K}(H),$$

for all $n \geq 1$, almost surely. Thus the addition of the penalty term distinguishes L^K as the unique minimizer among all matrices of rank in $\{r, \ldots, K/4\}$, almost surely. We will now show that for n sufficiently large, L^K almost surely uniquely minimises $S_{n,K}$, even when considering matrices of rank in $\{1, \ldots, r-1\}$.

Let $\theta \in \mathbb{R}^{K\times K}$, and $\theta_q \in \mathbb{R}^{K^*\times K^*}$ be a matrix obtained by keeping only rows and columns of θ with indices in the index set q and define the function $\tau : [K^*] \to q$ such that $\theta_q(i,j) = \theta(\tau(i),\tau(j))$. Then, by the coupling construction of $\boldsymbol{t}_K$ and $\boldsymbol{t}_{K^*}$, it almost surely holds that

$$\begin{aligned}
M_K(\theta) &= \sum_{\substack{i,j\in[K]\\ |i-j|>\lceil K/4\rceil}} K^{-2}\Big|\theta(i,j) - r(t_{i,K},t_{j,K})]\Big|^2 \\
&= \sum_{\substack{i,j\in q\\ |\tau^{-1}(i)-\tau^{-1}(j)|>\lceil K^*/4\rceil}} K^{-2}\Big|\theta(i,j) - r(t_{i,K},t_{j,K})]\Big|^2 + \\
&\qquad \sum_{\substack{i,j\in q\\ |\tau^{-1}(i)-\tau^{-1}(j)|\leq\lceil K^*/4\rceil\\ |i-j|>\lceil K/4\rceil}} K^{-2}\Big|\theta(i,j) - r(t_{i,K},t_{j,K})]\Big|^2 + \\
&\qquad + \sum_{\substack{i,j\notin q\\ |i-j|>\lceil K/4\rceil}} K^{-2}\Big|\theta(i,j) - r(t_{i,K},t_{j,K})]\Big|^2 \\
&\geq \sum_{a,b\in[K^*]:|a-b|>\lceil K^*/4\rceil} K^{-2}\Big|\theta_q(a,b) - r(t_{a,K^*},t_{b,K^*})]\Big|^2 \\
&= \left(\frac{K^*}{K}\right)^2 M_{K^*}(\theta_q).
\end{aligned}$$

As θ ranges over $\mathbb{R}^{K\times K}$ matrices of rank at most $r-1$, the sub-matrix θ_q ranges over $\mathbb{R}^{K^*\times K^*}$ matrices of rank at most $r-1$. It follows that

$$\inf_{\substack{\theta\in\Theta_K\\ \mathrm{rank}(\theta)<r}} M_K(\theta) \geq \left(\frac{K^*}{K}\right)^2 \inf_{\substack{\theta\in\Theta_{K^*}\\ \mathrm{rank}(\theta)<r}} M_{K^*}(\theta) = \left(\frac{K^*}{K}\right)^2 \epsilon(\boldsymbol{t}_{K^*}), \quad \text{almost surely.}$$

Since $K^* = 4r+4$, $\epsilon(\boldsymbol{t}_{K^*}) > 0$ almost surely, as $\boldsymbol{t}_{K^*}$ is uniformly distributed on $\mathcal{T}_{K^*}$. Thus, if n is sufficiently large so that $\tau_n < (K^*)^2\epsilon(\boldsymbol{t}_{K^*})/r$ a.s., it holds that

$$\left(\frac{K^*}{K}\right)^2 \epsilon(\boldsymbol{t}_{K^*}) > r\tau_n K^{-2}, \qquad \text{almost surely.}$$

Note that the critical $N(\boldsymbol{t}_{K^*})$ that n needs to surpass for this to hold true does not depend on K but only on the distribution of $\boldsymbol{t}_{K^*}$, which is the uniform distribution on $\mathcal{T}_{K^*}$, and on the precise rate of convergence of τ_n to zero (which is for us to choose). Consequently, $N(\boldsymbol{t}_{K^*})$ *does not grow* with K. Summarising, provided that n is sufficiently large,

$$\inf_{\substack{\theta\in\Theta_K \\ \mathrm{rank}(\theta)<r}} M_K(\theta) > r\tau_n K^{-2} = S_{n,K}(L^K), \qquad \text{almost surely.}$$

But

$$\inf_{\substack{\theta\in\Theta_K \\ \mathrm{rank}(\theta)<r}} S_{n,K}(\theta) = \inf_{\substack{\theta\in\Theta_K \\ \mathrm{rank}(\theta)<r}} \left\{M_K(\theta) + \tau_n K^{-2}\mathrm{rank}(\theta)\right\} > \inf_{\substack{\theta\in\Theta_K \\ \mathrm{rank}(\theta)<r}} M_K(\theta), \text{almost surely.}$$

which shows L^K to be the *unique* minimizer of $S_{n,K}(\cdot)$ on Θ_K almost surely, for n sufficiently large. Rephrasing, we have shown that for almost any grid $\boldsymbol{t}_K \in \mathcal{T}_K$, there exists an N (independent of K), such for all $n \geq N$, the matrix L^K is the unique minimum of $S_{n,K}$.
From now on, take $n \geq N(\boldsymbol{t}_{K^*})$, and note that all statements are valid almost surely with respect to the random grid $\boldsymbol{t}_K$. Our strategy will be to make use of van der Vaart and Wellner [20, Thm 3.4.1]. To this aim, let $\gamma > 0$, consider the difference

$$\Delta(\theta) = S_{n,K}(\theta) - S_{n,K}(L^K) = M_K(\theta) - M_K(L^K) + \tau_n\mathrm{rank}(\theta)K^{-2} - r\tau_n K^{-2},$$

and note that for some $p \in [0,1]$ and $\theta_* = pL^K + (1-p)\theta$, we have

$$\Delta(\theta) = \langle M_K'(L^K), (\theta - L^K)\rangle_{\mathrm{F}} + \frac{1}{2}\langle M_K''(\theta_*)(\theta - L^K), (\theta - L^K)\rangle_{\mathrm{F}} + \tau_n K^{-2}(\mathrm{rank}(\theta) - r),$$

where we have slightly abused notation, viewing the matrices as vectors. But $M_K'(\theta) = 2K^{-2}P^K \circ (\theta - R^K)$ and $M_K''(\tilde{\theta})\theta = 2K^{-2}P^K \circ \theta$. Noticing that $M_K'(L^K) = 0$ (by Theorem 2), we have

$$\begin{aligned}
\Delta(\theta) &= K^{-2}\langle P^K \circ (\theta - L^K), \theta - L^K\rangle_{\mathrm{F}} + \tau_n K^{-2}(\mathrm{rank}(\theta) - r) \\
&= K^{-2}\langle P^K \circ (\theta - L^K), P^K \circ (\theta - L^K)\rangle_F + \tau_n K^{-2}(\mathrm{rank}(\theta) - r) \\
&= K^{-2}\|P^K \circ (\theta - L^K)\|_{\mathrm{F}}^2 + \tau_n K^{-2}(\mathrm{rank}(\theta) - r) \\
&\leq K^{-2}\|\theta - L^K\|_{\mathrm{F}}^2 + \tau_n K^{-2}|\mathrm{rank}(\theta) - r|.
\end{aligned}$$

It follows that $\sup_{\theta\in\Theta_K : d_n(L^K,\theta)<\gamma} |\Delta(\theta)| \leq \gamma^2$. Next, consider:

$$\begin{aligned}
D(\theta) &= \mathbb{S}_{n,K}(\theta) - S_{n,K}(\theta) - \mathbb{S}_{n,K}(L^K) + S_{n,K}(L^K) \\
&= \mathbb{M}_{n,K}(\theta) - M_K(\theta) - \mathbb{M}_{n,K}(L^K) + M_K(L^K).
\end{aligned}$$

We expand $(\mathbb{M}_{n,K} - M_K)$ in a first order Taylor expansion with Lagrange remainder, around L^K, which gives for a certain $\tilde{p} \in [0,1]$ and $\tilde{\theta} = \tilde{p}L^K + (1-\tilde{p})\theta$:

$$\begin{aligned}
D(\theta) &= \langle \mathbb{M}_{n,K}'(\tilde{\theta}), \theta - L^K\rangle_{\mathrm{F}} - \langle M_K'(\tilde{\theta}), \theta - L^K\rangle_{\mathrm{F}} \\
&= K^{-2}\langle 2P^K \circ (\tilde{\theta} - R_n^K), \theta - L^K\rangle_{\mathrm{F}} - K^{-2}\langle 2P^K \circ (\tilde{\theta} - R^K), (\theta - L^K)\rangle_{\mathrm{F}} \\
&= K^{-2}\langle 2P^K \circ \tilde{\theta} - 2P^K \circ \tilde{\theta} - 2P^K \circ R_n^K + 2P^K \circ R^K, \theta - L^K\rangle_{\mathrm{F}} \\
&\leq K^{-2}\|2P^K \circ (R_n^K - R^K)\|_{\mathrm{F}}\|\theta - L^K\|_{\mathrm{F}} \leq 2K^{-1}\|R_n^K - R^K\|_{\mathrm{F}}K^{-1}\|\theta - L^K\|_{\mathrm{F}}.
\end{aligned}$$

Since $\mathbb{E}\|X\|_{L^2}^4 < \infty$, the process $X(s)X(t)$ is trace class on $[0,1]^2$ and thus has a continuous covariance kernel on $[0,1]^4$ (and consequently a continuous variance function on $[0,1]^2$). Assume without loss of generality

that $\mathbb{E}X = 0$. Since the observations $X_i(t_j)$ are independent for distinct i, and since $X_m(t_j)X_m(t_j)$ is an unbiased estimator of $\mathbb{E}[X(t_j)X(t_j)]$, we have

$$\begin{aligned} K^{-2}\mathbb{E}\|R_n^K - R^K\|_F^2 &= \sum_{i=1}^K\sum_{j=1}^K K^{-2}\mathbb{E}\left[\frac{1}{n}\sum_{m=1}^n X_m(t_{i,K})X_m(t_{j,K}) - \mathbb{E}\left[X(t_{i,K})X(t_{j,K})\right]\right]^2 \\ &= \frac{K^{-2}}{n}\sum_{i=1}^K\sum_{j=1}^K \mathrm{Var}[X(t_{i,K})X(t_{j,K})] \\ &\leq \frac{1}{n}\sup_{(s,t)\in[0,1]^2}\mathrm{Var}[X(s)X(t)] = \frac{C}{n}, \end{aligned}$$

and $C = \sup_{[0,1]^2}\mathrm{Var}[X(s)X(t)] < \infty$. Note that C is uniform in K. In summary we may conclude that

$$\mathbb{E}\left\{\sup_{\theta\in\Theta_K : d_n(\theta, L^K)<\gamma}|D(\theta)|\right\} \leq 2\gamma K^{-1}\mathbb{E}\|R_n^K - R^K\|_{\mathrm{F}} \leq 2\gamma\sqrt{K^{-2}\mathbb{E}\|R_n^K - R^K\|_F^2} = 2\gamma\sqrt{\frac{C}{n}}.$$

It now follows from van der Vaart and Wellner [20, Thm 3.4.1], that if $\hat{L}_n^K$ is an approximate maximiser of $\mathbb{S}_{n,K}$, in the sense given by the assumptions, then it holds that

$$n\, d_n^2(\hat{L}_n^K, L^K) = nK^{-2}\|\hat{L}_n^K - L^K\|_F^2 + n\left(\tau_n K^{-2}|\mathrm{rank}(\hat{L}_n^K) - r|\right) = O_{\mathbb{P}}(1),$$

where the $O_{\mathbb{P}}(1)$ term does not depend on K. We may conclude that

$$n\|\hat{\mathscr{L}}_n^K - \mathscr{L}^K\|_{\mathrm{HS}}^2 = nK^{-2}\|\hat{L}_n^K - L^K\|_F^2 = O_{\mathbb{P}}(1),$$

$$n\tau_n K^{-2}|\mathrm{rank}(\hat{\mathscr{L}}_n^K) - r| = O_{\mathbb{P}}(1),$$

where $O_{\mathbb{P}}(1)$ is uniform in K. We now turn our attention to the estimated eigenfunctions. Since these are finitely many, we will omit the index indicating the order of an eigenfunction for tidiness, and consider an eigenfunction η. Let η^K be the K-resolution step function approximation of η,

$$\eta^K(x) = \sum_{j=1}^K \eta(t_{j,K})\mathbf{1}\{x \in I_{j,K}\}.$$

Then, by Taylor expanding,

$$\int_0^1 \left(\eta(x) - \eta^K(x)\right)^2 dx = \sum_{j=1}^K \int_{I_{j,K}} \left(\eta(x) - \eta(t_{j,K})\right)^2 dx \leq \sum_{j=1}^K K^{-2}\|\eta'\|_\infty^2 = \frac{\|\eta'\|_\infty^2}{K}.$$

It follows that

$$\begin{aligned} \|\hat{\eta} - \eta\|_{L^2}^2 &\leq 2\left\|\hat{\eta} - \eta^K\right\|_{L^2}^2 + 2\left\|\eta^K - \eta\right\|_{L^2}^2 \\ &\leq c\|\hat{\mathscr{L}}_n^K - \mathscr{L}^K\|_{\mathrm{HS}}^2 + \frac{2\|\eta'\|_\infty^2}{K} = O_{\mathbb{P}}(n^{-1}) + \frac{2\|\eta'\|_\infty^2}{K}, \end{aligned}$$

and the $O_{\mathbb{P}}(n^{-1})$ has been shown to be uniform in K. The constant c can be chosen uniformly over the order of eigenfunction, since there are only $r < \infty$ eigenfunctions to consider. The convergence rate for $\sup_j|\hat{\lambda}_j - \lambda_j|$ follows from the inequality $\sup_j|\hat{\lambda}_j - \lambda_j| \leq \|\hat{\mathscr{L}}_n^K - \mathscr{L}\|_{\mathrm{HS}}$ (see, e.g., Bosq [1, Eq. 4.43]). □

Proof of Corollary 1. We have

$$\left\|\hat{\mathscr{B}}_n^K - \mathscr{B}\right\|_{\mathrm{HS}} \leq \left\|\hat{\mathscr{B}}_n^K - \mathscr{B}^K\right\|_{\mathrm{HS}} + \left\|\mathscr{B}^K - \mathscr{B}\right\|_{\mathrm{HS}}$$

$$= K^{-1}\left\|\hat{B}_n^K - B^K\right\|_{\mathrm{F}} + \left\|\mathscr{B}^K - \mathscr{B}\right\|_{\mathrm{HS}}, \tag{9.2}$$

where the second term on the right hand side converge in probability to zero as K goes to infinity. Let C be the set of banded $K \times K$ matrices of bandwidth at most $\lceil K/4 \rceil$ and D the set of non-negative definite $K \times K$ matrices. Let P_C and P_D be the projections onto C and D respectively and define $P = \lim_{l\to\infty}(P_C P_D)^l$. The first term on the right hand side of equality (9.2) can be written as

$$\begin{aligned} K^{-1}\left\|\hat{B}_n^K - B^K\right\|_{\mathrm{F}} &= K^{-1}\left\|P(R_n^K - \hat{L}_n^K) - P(R^K - L^K)\right\|_{\mathrm{F}} \\ &\leq K^{-1}\left\|(R_n^K - \hat{L}_n^K) - (R^K - L^K)\right\|_{\mathrm{F}}, \end{aligned}$$

where from the consistency of $\hat{L}_n^K$ and the definition of R_n^K we obtain that the right hand side term of the last inequality converge to zero as n goes to infinity. The consistency of the eigenfunctions $\hat{\psi}_m$ follows directly from the consistency of $\hat{\mathscr{B}}_n^K$. This completes the first part of the theorem. Now assume that b is continuously differentiable on the δ-band. In what follows, we will write $I_{ij} = I_{i,K} \times I_{j,K}$ for tidiness. We start with the decomposition:

$$\left\|\hat{\mathscr{B}}_n^K - \mathscr{B}\right\|_{\mathrm{HS}}^2 \leq 2\left\|\hat{\mathscr{B}}_n^K - \mathscr{B}^K\right\|_{\mathrm{HS}}^2 + 2\left\|\mathscr{B}^K - \mathscr{B}\right\|_{\mathrm{HS}}^2.$$

Since the kernel b is C^1 inside the δ-band and zero outside it, we write,

$$\begin{aligned} \int_0^1\int_0^1 (b(x,y) - b_K(x,y))^2 dxdy &= \iint_A (b(x,y) - b_K(x,y))^2 dxdy \\ &\quad + \iint_{A^c} (b(x,y) - b_K(x,y))^2 dxdy, \end{aligned}$$

where $A = \bigcup_{i,j}\left\{I_{ij} : \{|x-y| = \delta\} \cap I_{ij} \neq \emptyset\right\}$ is the union of rectangles I_{ij} that intersect with the boundary of the δ-band. On this set, $(b - b_K)^2$ is bounded by $\sup_{x,y,z,w}|b(x,y) - b_K(z,w)|^2 = C < \infty$, and so $\iint_A (b(x,y) - b_K(x,y))^2 dxdy$ is bounded by $|A|K^{-2}C$, where $|A| = \#\{(i,j) : I_{ij} \subset A\}$. Noting that $|A|$ is of the order of K, we have

$$\iint_A (b(x,y) - b_K(x,y))^2 dxdy = O(K^{-1}).$$

For the other term, we decompose the integral and use Taylor's theorem to see that

$$\begin{aligned} \iint_{A^c} (b(x,y) - b_K(x,y))^2 dxdy &= \iint_{\{|x-y|\leq\delta\}\cap A^c} (b(x,y) - b_K(x,y))^2 dxdy \\ &\quad + \iint_{\{|x-y|\geq\delta\}\cap A^c} (b(x,y) - b_K(x,y))^2 dxdy \end{aligned}$$

$$\begin{aligned} &= \iint_{\{|x-y|\leq\delta\}\cap A^c} (b(x,y) - b_K(x,y))^2 dxdy \\ &\leq \sum_{i,j=1}^K \mathbf{1}\Big\{I_{ij} \subset \big\{\{|x-y| \leq \delta\} \cap A^c\big\}\Big\} \iint_{I_{ij}} 2K^{-2} \sup_{(x,y)\in I_{ij}} \|\nabla b(x,y)\|_2^2 dxdy \\ &\leq \nu(K) 2K^{-4} \sup_{|x-y|<\delta} \|\nabla b(x,y)\|_2^2 = O(K^{-3}), \end{aligned}$$

where $\nu(K)$ is the number of rectangles within the δ-band, and thus satisfies $\nu(K) = O(K)$. In conclusion,

$$\left\|\mathscr{B}^K - \mathscr{B}\right\|_{\mathrm{HS}}^2 = O(K^{-1}).$$

For the other term, we note that, almost everywhere on $\mathcal{T}_K$,

$$\|\hat{\mathscr{B}}_n^K - \mathscr{B}^K\|_{\mathrm{HS}}^2 \leq \|\mathscr{D}_n^K - \mathscr{B}^K\|_{\mathrm{HS}}^2 \leq 2\|\hat{\mathscr{L}}_n^K - \mathscr{L}^K\|_{\mathrm{HS}}^2 + 2\|\hat{\mathscr{R}}_n^K - \mathscr{R}^K\|_{\mathrm{HS}}^2 = O_{\mathbb{P}}(n^{-1}),$$

where $\mathscr{D}_n^K$ is the operator corresponding to the matrix $\Delta_n^K = R_n^K - \hat{L}_n^K$ and, with the $O_{\mathbb{P}}(n^{-1})$ being uniform in K (as has been shown in the proof of Theorem 3). The convergence rate for the estimated eigenfunctions is now obtained by the fact that $\|\hat{\psi}_j - \psi_j\|_{L^2} \leq 2\sqrt{2}\sigma_j^{-1}\|\hat{\mathscr{B}}_n^K - \mathscr{B}\|_{\mathrm{HS}}$ (Bosq [1, Lemma 4.3]).

To complete the proof, we need to show that if $b \in C^1([0,1]^2)$, then $\left\|\mathscr{B}^K - \mathscr{B}\right\|_{\mathrm{HS}}^2 = O(K^{-3})$. In this case we may Taylor expand to write

$$\begin{aligned}\int_0^1\int_0^1 (b(x,y) - b_K(x,y))^2 dxdy &\leq \sum_{i,j=1}^K \iint_{I_{ij}} 2K^{-2} \sup_{(x,y)\in I_{ij}} \|\nabla b(x,y)\|_2^2 dxdy \\ &= \sum_{i,j=1}^K 2K^{-4} \sup_{(x,y)\in I_{ij}} \|\nabla b(x,y)\|_2^2 = O(K^{-3}),\end{aligned}$$

since there are at most order K terms of the form $\sup_{(x,y)\in I_{ij}} \|\nabla b(x,y)\|_2^2$ that are non-zero, by the bandedness of the kernel b.

The convergence rates for $\sup_j |\hat{\beta}_j - \beta_j|$ follow from the inequality $\sup_j |\hat{\beta}_j - \beta_j| \leq \|\hat{\mathscr{B}}_n^K - \mathscr{B}\|_{\mathrm{HS}}$ (see, e.g., Bosq [1, Eq. 4.43]).

□

Proof of Corollary 3. Since $K \geq K^*$, it must be that

$$\|\hat{\mathscr{L}}_n^K - \mathscr{L}^K\|_{\mathrm{HS}} = O_{\mathbb{P}}(n^{-1/2}), \quad \& \quad \|\hat{\mathscr{B}}_n^K - \mathscr{B}^K\|_{\mathrm{HS}} = O_{\mathbb{P}}(n^{-1/2}),$$

almost everywhere on $\mathcal{T}_K$ (as has been shown in the proof of Theorem 3 and of Corollary 1). Consequently, for almost all grids in $\mathcal{T}_K$,

$$\|\hat{\mathscr{R}}_n^K - \mathscr{R}^K\|_{\mathrm{HS}} \leq \|\hat{\mathscr{L}}_n^K - \mathscr{L}^K\|_{\mathrm{HS}} + \|\hat{\mathscr{B}}_n^K - \mathscr{B}^K\|_{\mathrm{HS}} = O_{\mathbb{P}}(n^{-1/2}).$$

It thus holds true that, for almost all grids in $\mathcal{T}_K$,

$$|\hat{\theta}_i - \theta_i^K| = O_{\mathbb{P}}(n^{-1/2}), \quad i = 1, \ldots, K,$$

where $\hat{\theta}_i$ (resp. θ_i^K) is the ith eigenvalue of $\hat{\mathscr{R}}_n^K$ (resp. $\mathscr{R}^K$). Since $\mathrm{rank}(\mathscr{R}^K) = K$, it must be that $\theta_1^K, \ldots, \theta_K^K > 0$. Letting $g(x) = x^{-1}\mathbf{1}\{x > 0\}$, and noting that g is differentiable at $\{\theta_i^K\}_{i=1}^K$, the delta method thus implies

$$\left|\frac{\mathbf{1}\{\hat{\theta}_i > 0\}}{\hat{\theta}_i} - \frac{\mathbf{1}\{\theta_i^K > 0\}}{\theta_i^K}\right| = \left|\frac{\mathbf{1}\{\hat{\theta}_i > 0\}}{\hat{\theta}_i} - \frac{1}{\theta_i^K}\right| = O_{\mathbb{P}}(n^{-1/2}), \quad i = 1, \ldots, K,$$

for almost all grids in $\mathcal{T}_K$. Now observe that

$$\hat{Y}_n^K := \hat{\Pi}_n(X^K) = \sum_{j=1}^{\hat{r}}\sum_{i=1}^K \mathbf{1}\{\hat{\theta}_i > 0\}\frac{\hat{\lambda}_j}{\hat{\theta}_i}\langle\hat{\varphi}_i, \hat{\eta}_j\rangle\langle\hat{\varphi}_i, X^K\rangle\hat{\eta}_j = \sum_{i=1}^K \mathbf{1}\{\hat{\theta}_i > 0\}\left(\frac{\langle\hat{\varphi}_i, X^K\rangle}{\hat{\theta}_i}\right)\hat{\mathscr{L}}_n^K\hat{\varphi}_i. \tag{9.3}$$

By the continuous mapping theorem, we know that the right hand side converges in probability to

$$\sum_{i=1}^K \left(\frac{\langle\varphi_i^K, X^K\rangle}{\theta_i^K}\right)\mathscr{L}^K\varphi_i^K = \sum_{j=1}^r\sum_{i=1}^K \frac{\lambda_j^K}{\theta_i^K}\langle\varphi_i^K, \eta_j^K\rangle\langle\varphi_i^K, X^K\rangle\eta_j^K = \Pi(X^K),$$

for almost all grids in $\mathcal{T}_K$, as $n \to \infty$. The fact that the rate of convergence is $O_{\mathbb{P}}(n^{-1/2})$ follows directly from the fact that each term of the summands in the right hand side of Equation 9.3 has been shown to converge at the rate $O_{\mathbb{P}}(n^{-1/2})$. The corresponding result follows for $\|\hat{W}_n^K - \Psi(X^K)\|_{L^2}$ by writing

$$\|\hat{W}_n^K - \Psi(X^K)\|_{L^2} = \|X^K - \hat{Y}_n^K - (X^K - \Pi(X^K))\|_{L^2} = \|\hat{Y}_n^K - \Pi(X^K)\|_{L^2}.$$

□

Corollary 2 follows directly from Theorem 3 and Corollary 1.

10. Appendix

This Appendix is structured as follows. Section 10.1 discusses the distorting effects of traditional FDA analysis on data that are characterised by two scales of variation in depth. Section 10.2 gives counterexamples that demonstrate that the combination of analyticity/banding assumptions is sharp (more precisely, that without more assumptions on the banded component, analyticity of the smooth component is necessary). Section 10.3 demonstrates that the scree-plot approach described in the main article indeed yield the rank-penalised estimator. Section 10.4 illustrates our methodology by means of a data analysis. Finally, Section 10.5 contains additional simulation results, probing varying combinations of sample size n and resolution K.

10.1. More on the Effect of Smoothing and PCA

Recall that our setup is

$$\rho(s,t) = \underbrace{\sum_{j=1}^{r} \lambda_j \eta_j(s)\eta_j(t)}_{\ell(s,t)} + \underbrace{\sum_{j=1}^{\infty} \beta_j \psi_j(s)\psi_j(t)}_{b(s,t)},$$

where: (1) $b(s,t) = 0$ for $|s-t| > \delta$, $0 < \delta < 1$; (2) $r < \infty$; (3) the $\{\eta_j\}$ are sufficiently smooth. The covariance kernel $\rho(\cdot,\cdot)$ of X admits its own uniformly convergent Mercer expansion,

$$\rho(s,t) = \sum_{j=1}^{\infty} \theta_j \varphi_j(s)\varphi_j(t) = \sum_{j=1}^{r} \lambda_j \eta_j(s)\eta_j(t) + \sum_{j=1}^{\infty} \beta_j \psi_j(s)\psi_j(t).$$

The question now is: what is the relationship between the system $\{\varphi_j\}$ and the systems $\{\eta_j\}$ and $\{\psi_j\}$? If it so happens that the $\{\eta_j\}_{j=1}^r$ system is orthogonal to the $\{\psi_j\}_{j=1}^\infty$ system, and we are fortunate enough that $\max_i \beta_i < \min_i \lambda_i$ then $\{\varphi_j = \eta_j : j \leq r\}$ and $\{\varphi_j = \psi_j : j > r\}$, and a direct Karhunen-Loève analysis will perfectly recover the smooth and rough variations. All that is required is a good rule for estimating the "truncation point" r (see e.g. Yao et al. [22], or Panaretos et al. [16] for AIC-type criteria), and the first few components of the expansion will give the smooth variation, while the remaining ones will give the rough variation, just as is typically assumed in FDA). Of course, if $\max_i \beta_i > \min_i \lambda_i$, then a direct Karhunen-Loève expansion will still recover the correct principal components of variation, but their order will not distinguish the smooth from the rough components.

However, if the $\{\eta_j\}_{j=1}^r$ are *not* orthogonal to the $\{\psi_j\}_{j=1}^\infty$ (as may very well happen in practice) more severe distortions will arise: it may very well happen that neither η nor ψ will be eigenfunctions of $\mathscr{R}$, so that we cannot identify the carriers of smooth and rough variation from direct PCA. Assume, for example, that no pair $\{\eta_i, \psi_j\}$ is orthogonal. Then,

(a) If $\beta_1 > \lambda_k$ for some k, it is clear that the eigenfunctions $\{\varphi_j\}_{j\geq k}$ will be linear combinations of $\{\psi_j\}_{j\geq 1}$ and $\{\eta_j\}_{j\geq k}$. Thus, we will neither be able to recover the smooth components of variation beyond order k, nor the rough components: the extracted components of variation from order k onwards will be confounded versions of smooth and rough components of variation.

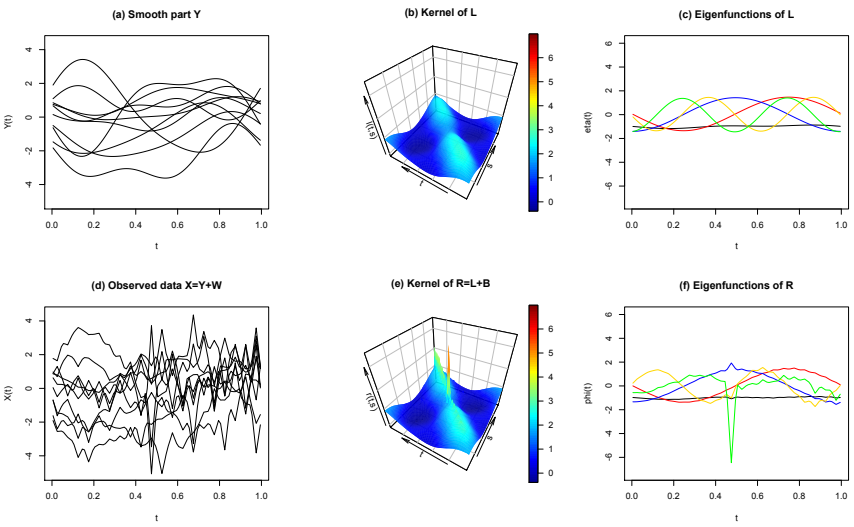

FIG 4. *Ten smooth curves $Y_1, \ldots, Y_{10}$ (plot (a)) corresponding to a rank 5 smooth operator $\mathscr{L}$ (plot (b)), with eigenfunctions η_j (plot (c), in decreasing order: black, red, blue yellow, green). To these smooth curves, we add uncorrelated banded rough processes, yielding observables $X_1, \ldots, X_{10}$ (plot (d)), whose covariance operator is $\mathscr{R}$ (plot (e)) with eigenfunctions φ_j (plot (f)).*

(b) Even if $\max_i \beta_i < \min_i \lambda_i$, it will still happen that $\phi_j \neq \psi_j$ for $j > r$ (since $\{\phi_j\}_{j>r}$ will be in the orthogonal complement of $\text{span}\{\eta_1, ..., \eta_r\}$, whereas $\{\psi_j\}_{j\geq 1}$ are not). In other words, the rough components of variation will be distorted (for example, if the $\{\psi_j\}_{j\geq 1}$ are locally supported, the $\{\phi_j\}_{j\geq r}$ will typically fail to be so). In fact, $\max_i \beta_i < \min_i \lambda_i$ alone does not even guarantee that $\phi_i = \eta_i$ for $i \leq r$. Depending on the spacings of $\{\lambda_j\}_{j=1}^r \cup \{\beta_j\}_{j\geq 1}$ it may happen that some of the $\{\phi_i\}_{i=1}^r$ could be linear combinations between the η_j and the ψ_j. Thus the smooth components of variation could be distorted too.

For instance, Figure 4 presents a simulated example where the data are constructed as the sum of a smooth process Y with trigonometric principal components and covariance $\mathscr{L}$ of rank 5; and a rough process W built as the sum of locally supported rough principal components with covariance $\mathscr{B}$ of (non-trivial) band $\delta = 0.05$ (see Section 8 for more details on this example; the eigenfunctions ψ_j are triangular functions locally supported on non-overlapping subintervals of length 0.05). The eigenvalues are chosen so that $\beta_1 > \lambda_5$. We see that $X = Y + W$ has fifth eigenfunction φ_5 (in green) that is a distorted version of η_5 (indeed a linear combination of η_5 and ψ_1). It is also clear that the eigenfunctions $\{\varphi_j\}_{j\geq 5}$ of X will typically not be locally supported (since they must be orthogonal to $\varphi_1, ..., \varphi_5$), in contrast to the true rough eigenfunctions $\{\psi_j\}_{j\geq 1}$ that were chosen to be locally supported. Finally, we see that even eigenfunctions of order lower than 5 have been affected (as we mentioned earlier this could happen too, depending on the spacings), and they contain artefacts resulting from confounding with rough eigenfunctions.

If smoothing were to also take place prior to a Karhunen-Loève expansion, then there could be a further confounding effect, at least for finite samples. Whether using splines or the PACE algorithm, we would essentially be convolving the discrete data with a kernel of some positive bandwidth $h > 0$ (spline smoothing can be seen as approximately kernel smoothing with an equivalent kernel, Silverman [19]). If the size of this bandwidth is comparable with δ (which it may be in finite samples), then the variations of scale δ due to the W component would propagate to larger scales, entangling the covariance of Y with that of W. Smoothing could also yield smoothed versions of the η_j and the ψ_j that are even further away from being orthogonal than initially (with the effects discussed earlier). The effect of smoothing is hard to quantify precisely, since the behaviour of the h parameter is typically understood asymptotically, and is usually chosen in a data dependent manner in finite samples (which can also be a source of further trouble, see for instance Opsomer et al. [15]).

If we could take the scale of W to be $\delta \approx 0$, then W would correspond to a generalised noise process. For instance, take the rough component W as being precisely white noise of level σ^2 (corresponding to taking $\mathscr{B}$ as being σ^2 times the identity), and interpret the equality $X(t) = Y(t) + W(t)$ in the weak sense $\langle X, f\rangle = \langle Y, f\rangle + \sigma^2 \int_0^1 f(t)dB_t$, for any $f \in L^2[0,1]$, and for $\{B_t\}$ a standard Brownian motion. In this case there is no confounding problem: the eigenfunctions φ_j corresponding to X would be exactly equal to the eigenfunctions η_j corresponding to Y, for all $j = 1, \ldots, r$ (the remaining φ_j could be taken to be any ONS for the orthogonal complement of $\text{span}\{\eta_1, \ldots, \eta_n\}$). Furthermore, the θ_j would simply satisfy $\theta_j = \lambda_j \mathbf{1}\{j \leq r\} + \sigma^2$. In particular their order would not change. Thus, smoothing (either by spline smoothing or by the PACE algorithm) followed by PCA would have essentially no distorting effects on our understanding of the covariation properties of X.

10.2. Analyticity and Uniqueness

In Remark 1 following Theorem 1 (conditions ensuring uniqueness of the decomposition $\mathscr{R} = \mathscr{L} + \mathscr{B}$), it was pointed out that the conditions of the theorem can actually be strictly weakened, while retaining the same conclusion. One can retain the bandedness assumption on $(\mathscr{B}_1, \mathscr{B}_2)$, but replace the assumption of requiring finite ranks and analytic eigenfunctions for $(\mathscr{L}_1, \mathscr{L}_2)$, by the weaker assumption that the kernels of $(\mathscr{L}_1, \mathscr{L}_2)$ be analytic on an open set $U \subset [0,1]^2$ that contains the larger of the two bands, $U \supset \{(s,t) \in [0,1]^2 : |t-s| \leq \max(\delta_1, \delta_2)\}$. However, this assumption cannot be further weakened, unless we are willing to make stronger assumptions on $\mathscr{B}$. If we seek *completely non-parametric* conditions for *unique recovery* of a decomposition $\mathscr{R} = \mathscr{L} + \mathscr{B}$ from knowledge of the sum $\mathscr{R}$, our assumptions are quite sharp: one cannot weaken one of them without strengthening the other. For instance, if we do not impose further restrictions on $\mathscr{B}$ than just bandedness, then analyticity of $\mathscr{L}$ is *necessary* and cannot be weakened. We now construct two counterexamples to demonstrate this.

10.2.1. Counterexample 1

Here we provide a counterexample to show that the analyticity assumption cannot be further weakened. Consider two stationary kernels ℓ_1 and ℓ_2 on $[0,10]^2$ defined as:

$$\ell_1(s,t) = \begin{cases} \dfrac{1}{1+(s-t)^2} + \exp\left\{-\dfrac{1}{1-(s-t)^2}\right\} & \text{if } |s-t| \leq 1 \\ \dfrac{1}{1+(s-t)^2} & \text{if } |s-t| > 1, \end{cases}$$

and

$$\ell_2(s,t) = \frac{1}{1+(s-t)^2}.$$

Note that: (1) ℓ_2 is analytic; (2) ℓ_1 is analytic, except on the line $|s-t| = 1$, and is C^∞ everywhere. Consequently, even though

$$\ell_1(s,t) = \ell_2(s,t), \qquad \forall |s-t| > 1,$$

it still happens that

$$\ell_1(s,t) \neq \ell_2(s,t), \qquad \forall |s-t| \leq 1.$$

We plot ℓ_1 and ℓ_2 as a function of $x = s - t$ in Figure 10.2.1 to illustrate the issue. Now define banded kernels, with a bandwidth of at most 1,

$$b_1(s,t) = 0,$$

and

$$b_2(s,t) = \begin{cases} \exp\left\{-\frac{1}{1-(s-t)^2}\right\} & \text{if } |s-t| \leq 1, \\ 0 & \text{if } |s-t| > 1. \end{cases}$$

We now have

$$\ell_1 + b_1 = \ell_2 + b_2,$$

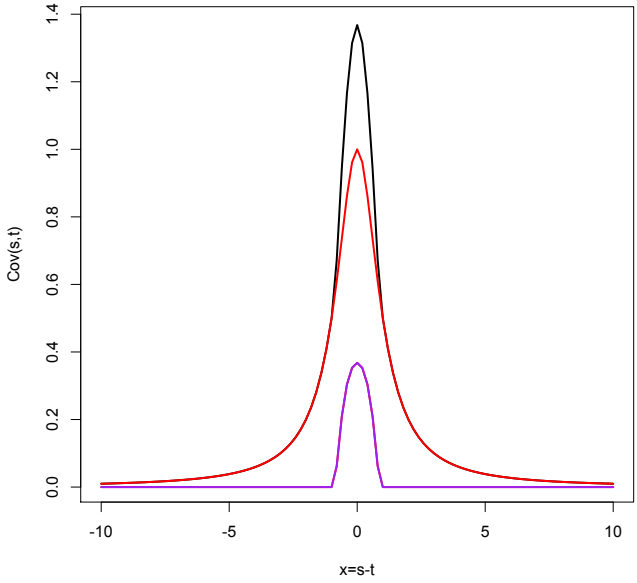

FIG 5. *The kernels ℓ_1 (in black), ℓ_2 (in red), and b_2 (in purple) as a function of $x = s - t$.*

but of course $\ell_1 \neq \ell_2$ and $b_1 \neq b_2$. Figure 10.2.1 shows ℓ_1 in black, ℓ_2 in red, and b_2 in purple, as a function of $x = s - t$.

10.2.2. Counterexample 2

The first counterexample included stationary kernels of infinite rank. We now show that analyticity remains a necessary assumption even in a finite rank situation. For some $\delta \in (0, 1)$ let

$$\varphi_\delta(x) = \begin{cases} \exp\left\{-\dfrac{1}{1-(x/\delta)^2}\right\} & \text{if } x \in [0, \delta), \\ 0 & \text{otherwise,} \end{cases} \tag{10.1}$$

and let $\psi(x)$ be an analytic function on $[0, 1]$ (for example $\psi(x) = x$). Define the covariance kernel

$$\rho(x, y) = \psi(x)\psi(y) + \phi_\delta(x)\phi_\delta(y),$$

and note that it has rank 2, while each of its summands has rank 1. Moreover, the component $\phi_\delta(x)\phi_\delta(y)$ is supported on $[0, \delta]^2$, and thus it is banded with bandwidth δ. It follows that we may define:

$$\begin{aligned} \ell_1(x, y) &= \psi(x)\psi(y), \\ \ell_2(x, y) &= \psi(x)\psi(y) + \phi_\delta(x)\phi_\delta(y), \\ b_1(x, y) &= \phi_\delta(x)\phi_\delta(y), \\ b_2(x, y) &= 0, \end{aligned}$$

such that $\ell_1 \neq \ell_2$ and $b_1 \neq b_2$ but

$$\ell_1 + b_1 = \ell_2 + b_2.$$

Note that once again the reason uniqueness fails is that analyticity does not hold on an open interval containing the band $\{|x-y|<\delta\}$: the kernel $\phi_\delta(x)\phi_\delta(y)$ is analytic on open neighbourhoods of any pair of points on the band $|x-y|=\delta$, except for two such points: the points $\{x=0,y=\delta\}$ and $\{x=\delta,y=0\}$.

We conclude this counterexample by noting that the fact that $\phi_\delta(x)\phi_\delta(y)$ was block-diagonal and of rank 1 is not essential: one can define the continuous superposition

$$\varphi(x,y)=\frac{1}{1-2\delta}\int_\delta^{1-\delta}\phi_u(x)\phi_u(y)du,$$

that will be supported on the entire band $\{|x-y|<\delta\}$ and will be of inifinite rank, and still repeat the same example by replacing $\phi_\delta(x)\phi_\delta(y)$ by $\varphi(x,y)$.

10.2.3. Discussion of the Counterexamples

The two counterexamples illustrate the source of the difficulty, and indicate how yet more counterexamples could be constructed. Let $\mathscr{G}$ be a smooth covariance, and $\mathscr{B}_1$ and $\mathscr{B}_2$ be some banded covariances (not even necessarily of the same bandwidth). Define $\mathscr{R}=\mathscr{G}+\mathscr{B}_1+\mathscr{B}_1$. Then, note that we can write:

$$\mathscr{R}=\underbrace{\mathscr{G}}_{\mathscr{L}}+\underbrace{\mathscr{B}_1+\mathscr{B}_1}_{\mathscr{B}}\qquad\text{or}\qquad\mathscr{R}=\underbrace{\mathscr{G}+\mathscr{B}_1}_{\mathscr{L}}+\underbrace{\mathscr{B}_2}_{\mathscr{B}}\qquad\text{or}\qquad\mathscr{R}=\underbrace{\mathscr{G}+\mathscr{B}_2}_{\mathscr{L}}+\underbrace{\mathscr{B}_1}_{\mathscr{B}}.$$

In particular, one can devise such decompositions for any combination of C^k assumptions imposed on $\mathscr{G}$, $\mathscr{B}_1$, and $\mathscr{B}_2$. It follows that the assumptions on analyticity/banding should be seen as describing *what is feasible* in a purely non-parametric setup. From that perspective, the assumptions are quite intuitive: if we want to separate two components Y and W that represent two different scales of variation, then W should have variations at most of some scale δ, and Y should have variations at a scale that is at least δ.

10.3. On the Scree Plot Approach for the Choice of Tuning Parameter

The aim of this section is to illustrate the correspondence between steps (B) and (B') in Section 7 of the main paper. Specifically, we will show how selecting a value $c>0$ and solving the problem

$$\min_{\theta\in\mathbb{R}^{K\times K}}\operatorname{rank}(\theta)\qquad\text{subject to }\left\|P^K\circ(R_n^K-\theta)\right\|_F^2<c,\tag{10.2}$$

corresponds to selecting a value $\tau>0$ and solving the problem

$$\min_{\theta\in\mathbb{R}^{K\times K}}\left\{\left\|P^K\circ(R_n^K-\theta)\right\|_F^2+\tau\operatorname{rank}(\theta)\right\}.\tag{10.3}$$

To do this, we first introduce some definitions and make some observations. Let

$$f(i)=\min_{\theta\in\mathbb{R}^{K\times K},\operatorname{rank}(\theta)\le i}\{\left\|P^K\circ(R_n^K-\theta)\right\|_F^2\},\qquad i=1,\dots,K,\tag{10.4}$$

be the fit at rank i, and extend f to the positive reals by linear interpolation. Call the graph of $u\mapsto f(u)$ the "scree plot". Observe that $f(u)$ is non-increasing. Without loss of generality, assume that $f(1)=1$ and $f(K)=0$, otherwise renormalise appropriately. Define f^{-1} to be

$$f^{-1}(c)=\inf\{x\in\mathbb{R}:f(x)\le c\}.$$

With these definitions in place, note that solving 10.2 for a given $c>0$ is equivalent to solving

$$\min_{\theta\in\mathbb{R}^{K\times K}}\left\|P^K\circ(R_n^K-\theta)\right\|_F^2\qquad\text{subject to }\operatorname{rank}(\theta)\le\lceil f^{-1}(c)\rceil,\tag{10.5}$$

Finally, define the *increments* of the scree plot as

$$\Delta_i:=f(i)-f(i+1)\ge0,\quad i=1,\dots,K-1;\quad\Delta(K):=0.$$

We now have

Lemma 2. *If $x \mapsto f(x)$ is strictly convex, then, for any constant $c > 0$, the problem 10.2 with constraint parameter c is equivalent to 10.3 with a tuning parameter in the range*

$$\max\{\Delta_j : j \geq \lceil f^{-1}(c) \rceil\} < \tau < \min\{\Delta_j : j \leq \lceil f^{-1}(c) \rceil - 1\}.$$

Furthermore, τ can be made arbitrarily small by choosing c to be arbitrarily small.

Proof. Choose $c > 0$ and let $q = \lceil f^{-1}(c) \rceil$. If we can choose a value of τ that simultaneously satisfies

$$\begin{aligned} \tau(q-j) + f(q-j) &> \tau q + f(q), \quad \forall\, j < q \\ \tau(q+j) + f(q+j) &> \tau q + f(q), \quad \forall\, j \geq 1 \end{aligned}$$

then a candidate matrix θ will be a solution to the penalised optimisation problem 10.3 with tuning parameter τ if and only if $\text{rank}(\theta) = q$ and $\|P^K \circ (R_n^K - \theta)\|_F^2 = f(q)$. In other words, the optima of the penalised problem 10.3 will coincide with the optima of the constrained problem 10.2.

We now examine when choosing such a τ is feasible. Notice that the two conditions that τ must satisfy are equivalent to:

$$\tau < \frac{f(q-j) - f(q)}{j}, \quad \forall\, j < q \qquad \& \qquad \tau > \frac{f(q) - f(q+j)}{j}, \quad \forall\, j \geq 1.$$

And so, by telescoping,

$$\frac{f(q-j) - f(q)}{j} = \underbrace{\frac{f(q-j) - f(q-j+1)}{j} + \ldots + \frac{f(q-1) - f(q)}{j}}_{j \text{ terms}},$$

and

$$\frac{f(q) - f(q+j)}{j} = \underbrace{\frac{f(q) - f(q+1)}{j} + \ldots + \frac{f(q+j-1) - f(q+j)}{j}}_{j \text{ terms}}.$$

We may thus re-write the conditions on τ as

$$\begin{aligned} \tau &< \frac{f(q-j) - f(q-j+1)}{j} + \ldots + \frac{f(q-1) - f(q)}{j}, \quad \forall\, j < q, \\ \tau &> \frac{f(q) - f(q+1)}{j} + \ldots + \frac{f(q+j-1) - f(q+j)}{j}, \quad \forall\, j \geq 1. \end{aligned}$$

By convexity of arithmetic averaging, a sufficient condition for the above to be true is to require

$$\begin{aligned} \tau &< f(i) - f(i+1) := \Delta_i, \quad \forall\, i \leq q-1, \\ \tau &> f(i) - f(i+1) = \Delta_i, \quad \forall\, i \geq q. \end{aligned}$$

Since $x \mapsto f(x)$ is strictly convex, the sequence Δ_i is strictly decreasing in i. It follows that the last two conditions are compatible, and we may choose any τ in the range

$$\max\{\Delta_j : j \geq \lceil f^{-1}(c) \rceil\} < \tau < \min\{\Delta_j : j \leq \lceil f^{-1}(c) \rceil - 1\},$$

while retaining the same optima for the two problems. Furthermore, since Δ_j can be made arbitrarily small for $j \leq \lceil f^{-1}(c) \rceil$ by choosing c to be sufficiently small, we see that τ can be taken to be arbitrarily small by appropriate choice of c.

□

Note that if $x \mapsto f(x)$ is convex, then it will almost surely be strictly convex since $\{f(i)\}_{i \geq 1}$ are continuous random variables. We conclude this section by establishing the validity of the elbow selection rule as sample size diverges.

Lemma 3. *Assume the same conditions and context as in Proposition 2. Then, and for almost all grids in $\mathcal{T}_K$, it holds that*

$$\limsup_{n\to\infty} f(i) = 0 \qquad \textit{almost surely,}$$

for all $i \geq r$ whereas

$$\liminf_{n\to\infty} f(i) > \frac{1}{2}\sum_{j=i+1}^{r} \zeta_j^2 > 0 \qquad \textit{almost surely,}$$

for all $i < r$, whenever $r > 1$. Here $r = \operatorname{rank}(L^K)$ is the true rank of $\mathscr{L}$, and $\{\zeta_i\}_{i=1}^r$ are non-zero eigenvalues of the symmetric $K \times K$ matrix U^K, obtained by retaining the top-right and bottom-left $r \times r$ submatrices of L^K, and setting all other entries equal to zero.

Proof. We will write $f_n(i)$ instead of $f(i)$ in order to highlight the dependence on n. Let $\mathcal{A}_K \subseteq \mathcal{T}_K$ be the set of grids for which Proposition 2 is valid, and fix a grid $\mathbf{t}_K \in \mathcal{A}_K$. Note that this suffices for the purposes of the proof, since $\mathcal{A}_K$ is of full Lebesgue measure. Now, note that

$$f_n(r) \leq \left\|P^K \circ (R_n^K - L^K)\right\|_F^2 \xrightarrow{\text{a.s.}} \left\|P^K \circ (R^K - L^K)\right\|_F^2 = 0,$$

where $r = \operatorname{rank}(L^K)$. Consequently, $f_n(j) \leq f_n(r) \xrightarrow{\text{a.s.}} 0$ for all $j \geq r$, and obviously

$$\limsup_{n\to\infty} f_n(i) = 0 \qquad \text{almost surely,}$$

for all $i \geq r$. We now turn to the second assertion. We will consider the case $i = r-1$ (the remaining cases follow similarly). Write $\zeta = \zeta_r > 0$ for the smallest eigenvalue of U^K. First, note that this must be non-zero, since Theorem 2 implies that all $r \times r$ minors of L^K are of full rank r.

We will argue by contradiction: suppose that the event $\{f_n(r-1) < \zeta^2/2 \text{ infinitely often}\}$ has positive probability. It follows that there exists a sequence θ_k of rank $r-1$ random matrices and a subsequence $\{R_{n_k}^K\}$ of $\{R_n^K\}$ such that

$$\left\|P^K \circ (R_{n_k}^K - \theta_k)\right\|_F^2 = \left\|P^K \circ R_{n_k}^K - P^K \circ \theta_k\right\|_F^2 < \zeta^2/2, \qquad \forall\, k \geq 1,$$

with positive probability. On the other hand, we know that

$$\left\|P^K \circ R_{n_k}^K - P^K \circ R^K\right\|_F^2 \xrightarrow{\text{a.s.}} 0.$$

Consequently, since $P \circ (L^K - R^K) = 0$, it follows that for all k sufficiently large,

$$\left\|P^K \circ \theta_k - P^K \circ L^K\right\|_F^2 < \zeta^2/2 + \zeta^2/2 = \zeta^2,$$

with positive probability. Now let ϑ_k denote the symmetric matrix formed by retaining the bottom-left and top-right $r \times r$ minors of θ_k, and setting the remaining elements equal to zero. Since our assumptions entail that $K \geq K^* = 4r + 4$, we now have:

1. By Theorem 4, U^K is of rank r, and of course ϑ_k is of rank at most $r-1$, for all k, with probability 1.
2. The event $\left\|P^K \circ \theta_k - P^K \circ L^K\right\|_F < \zeta^2$ has positive probability, and thus the event $\|\vartheta_k - U^K\|_F^2 < \zeta^2$ also has positive probability.

These two conclusions constitute a contradiction: the closest element to U^K from within the set $\{\theta : \operatorname{rank}(\theta) = r-1\}$ is the $(r-1)$-spectral truncation of U^K, and this has squared Frobenius distance from U^K equal to ζ^2. This concludes the proof.

□

10.4. Data Analysis: Application to Air Pollution Data

As an illustration of our method, we analyse a data set related to the air quality in the city of Geneva, Switzerland. The data are comprised of measurements of the concentration of nitrogen dioxide (NO2) in the air (in micrograms per cubic meter), that have been recorded hourly at the "L'Ile" station, starting on the second Monday of September and until the second Sunday of November from 2005 to 2011. The data set can be accessed at:

`http://ge.ch/air/qualite-de-lair/requete-de-donnees`

Viewed as functional data, these measurements yield $n = 62$ curves corresponding to the different weeks, and each of these curves is evaluated at $K = 168$ points, corresponding to 7 days (from Monday to Sunday) times 24 hours. The raw curves and their empirical covariance function are plotted in Figure 6.

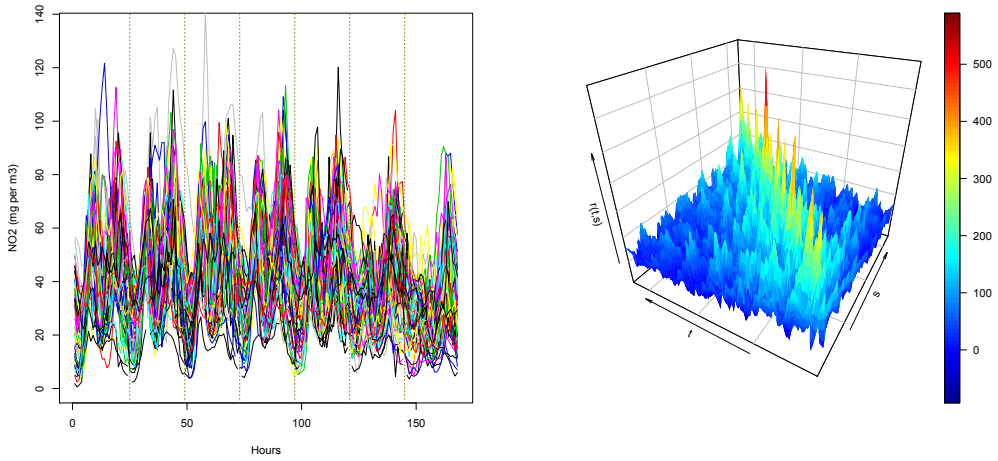

FIG 6. *The dataset is depicted on the left, and its empirical covariance function on the right.*

For these particular data, we are expecting the covariance kernel r to decompose into a component ℓ capturing variation at the time-scale of a week, and a second component b, capturing day-specific variation, thus essentially being concentrated around a band. The natural choice of upper bound for δ is thus 0.15, corresponding to removing a band of width $\delta \times K = 24$ hours in the discrete setup. In order to pick the rank r, we solved the optimisation problem 10.2 (as described in the main body) for $i = 1, \ldots, 7$, and we plotted the functions $f(i) = \|P^K \circ (R_n^K - \hat{C}_i \hat{C}_i^\top)\|_F^2$ and the ratio $r(j) = f(j)/f(j+1)$, for $j = 1, \ldots, 6$ on Figure 7. Our estimated rank should be the point i where the function f levels out, or equivalently, the point j for which the ratio r becomes a constant close to 1. The obvious choice was subsequently $\hat{r} = 3$ and our estimator of L_n^K is given by $\hat{L} = \hat{C}_3 \hat{C}_3^\top$.

A very slightly smoothed version of $\hat{L}$ is plotted on Figure 8, and the same figure plots its three corresponding eigenfunctions. These eigenfunctions represent variation that propagates globally throughout the whole week. The first eigenfunction appears to represent fluctuation of the overall level of concentration on a weakly basis – this upward/downward shift does have finer structure within each day, but: (a) these still represent fluctuations coupled/correlated during *all* mornings/afternoons in a week, and (b) the intraday structure of the eigenfunction reveals a morning and an afternoon peak of opposite sign, roughly reflecting that this mean level shift is purely weakly, and does not differ noticeably from day to day. The second eigenfunction appears to capture early/late week effects, showing that the period from Thursday to Sunday has a higher level of variation, which in fact correlates negatively with variation from Monday to Wednesday. Finally, the third eigenfunction seems to capture periodic day/night variation, as it propagates throughout the week, and it is clearly noticeable how this variation increases during the weekend.

The estimates of the covariance function b and of its first three eigenfunctions are plotted in Figure 9. A striking feature is that the eigenfunctions are almost exactly locally supported, though this was nowhere

enforced explicitly – they represent genuinely short scale variations that are uncorrelated across lengthier time scales. Each represents variation that is specific to a particular period in the week: the first chiefly during weekends, the second mostly during the early week, and the third more around mid-week (note that the corresponding eigenvalues are rather close in magnitude, so the order to the three eigenfunctions is not well-distinguished: these are effects of approximately equal magnitude).

These local fluctuations would have been annihilated by a traditional smooth plus PCA approach: Figure 10 depicts the six leading eigenfunctions of an estimate of l obtained by a Fourier basis smoothing with a roughness penalty approach (we use the Fourier rather than spline basis to respect the periodic nature of the data). The first three of these present overall features that not dissimilar to those given by our approach, albeit a bit more rough (this comes as no surprise, since the previous analysis shows that we are in a "well-ordered" scenario). But the next three eigenfunctions are supported globally and are completely uninterpretable. This is a consequence of the fact that they are constrained to be orthogonal to the first three (see the discussion in bullet point (b) of Section 10.1). To complicate matters further, the leading three eigenfunctions account for only 52% of the total variance, whereas the next three account for a further 16% – meaning that the one cannot rely on the first three eigenfunctions alone for their analysis, and needs all six to approach the traditional 80% threshold.

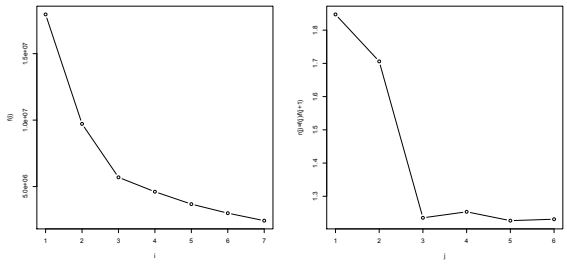

FIG 7. *Illustration of the scree plot approach to rank selection. On the left we plotted the function* $f(i) = \|P^K \circ (R_n^K - \hat{C}_i \hat{C}_i^\top)\|_F^2$ *for* $i = 1, \ldots, 7$*, and on the right the ratio* $r(j) = f(j)/f(j+1)$ *for* $j = 1, \ldots, 6$.

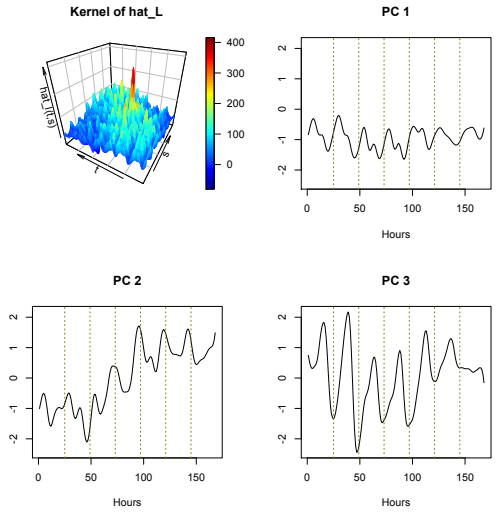

FIG 8. *Lightly smoothed estimate of the covariance function* ℓ *and the corresponding three eigenfunctions. Vertical dotted lines indicate the different days of the week, starting with Monday as the first block.*

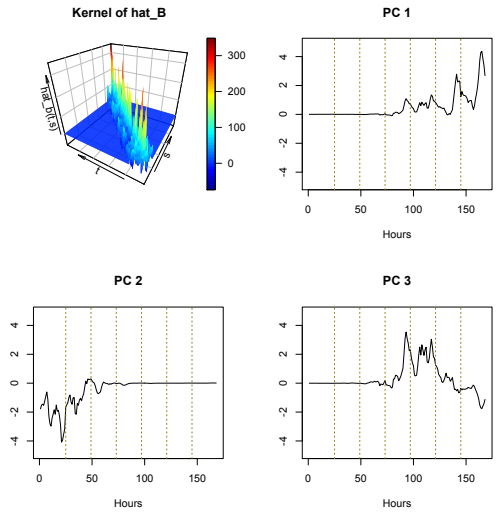

FIG 9. *Estimation of the covariance function b and of its first three eigenfunctions. The dotted lines indicate the different days of the week.*

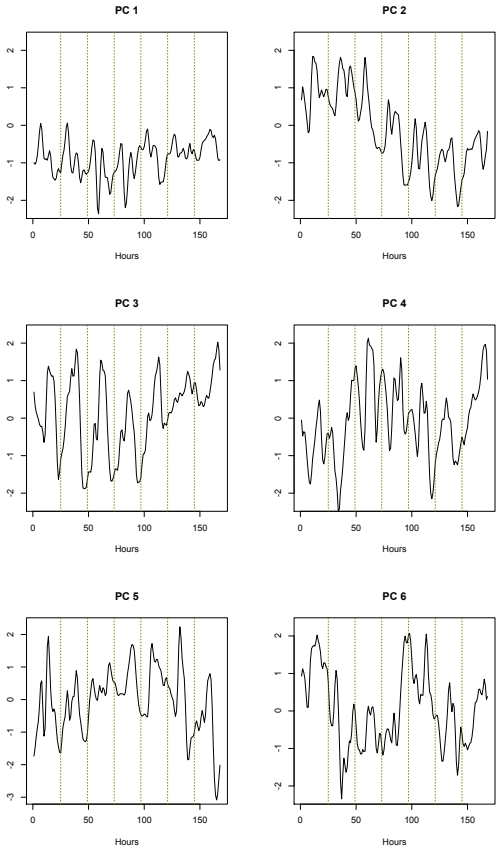

FIG 10. *The first six eigenfunctions of the estimate of l obtained by smoothing with a roughness penalty the empirical covariance matrix. The dotted lines indicate the different days of the week.*

10.5. Additional Simulation Results

This section contains additional plots from the simulation presented in Section 8 of the main article, as well as further simulation results. It was observed in Section 8.1 that one may slightly underestimate the rank when employing the scree-plot approach, especially when data are generated under regime 2 (interlaced eigenvalues). In order to appreciate the impact of rank misspecification, we have calculated the normalised errors $\mathrm{Err}(\cdot)$ of the estimators obtained with a rank choice of $2, 4$ and 5 when the true rank is 3 and with a rank choice of $3, 4, 6$ and 7 when the true rank is 5 for four different cases of scenario A (namely $\delta = 0.05$ and $\delta = 0.1$ in the interlaced and non-interlaced regimes); we used 100 replications for each case. Boxplots of the ratio between our method's error when the correct rank is used (in the denominator) and the error of our method when the rank is misspecified (in the numerator) are depicted in Figure 11. The red horizontal lines on the graphics indicate the level 1. It is clear that *underestimation* of the rank leads to more severe effects than *overestimation*. In particular, overestimation of the rank seems to not affect performance, except in isolated outlying cases. This explains our earlier recommendation that one shoud not hesitate to choose a larger rank when in doubt.

As mentioned in Section 8.2, we also ran additional simulations to study the performance of our method for different combinations of sample size n and grid points K. For the scenario A (FB + MA), rank/bandwidth combination (1–6), and the two regimes considered in the paper, we simulate 100 replications for the 6 different combinations of the sample size n and number of grid points K given in the Table 6.

n	300	300	300	100	100	100
K	25	50	150	25	50	100

TABLE 6
Different values of the number of curves and number of grid points.

For each simulation setup, we calculate the 100 normalised errors $\mathrm{Err}(u) = (||u - L_n^K||_F)/||L_n^K||_F$ for our method, the PACE method, the truncation of the Karhunen-Loève (KL) expansion method and the spline smoothing method. We then form the ratio between our method's error (in the denominator) and the error of each of the three other methods (in the numerator). The first quartiles, medians and third quartiles of the resulting distributions are presented in Tables 7 and 8. The medians exceeding 1 have been highlighted in bold. We see that our method continues to perform considerably better than the benchmark methods, regardless of the ratio of n/K in the case of Regime 1. In Regime 2, our method performs better or comparably to other methods in almost all combinations. The only exceptions are in the sparse regimes ($\{n = 100, K = 25\}$ and $\{n = 300, K = 25\}$): even in these cases, our method outperforms other methods when the rank is 3, but starts to underperform when the rank is 5. However, note when the rank is 5 and $K = 25$, we are the boundary of our identifiability theorem (which requires that $r \leq (1/2 - \delta)K - 1$), and of course the boundary itself applies to the population version, whereas here one is dealing with finite samples.

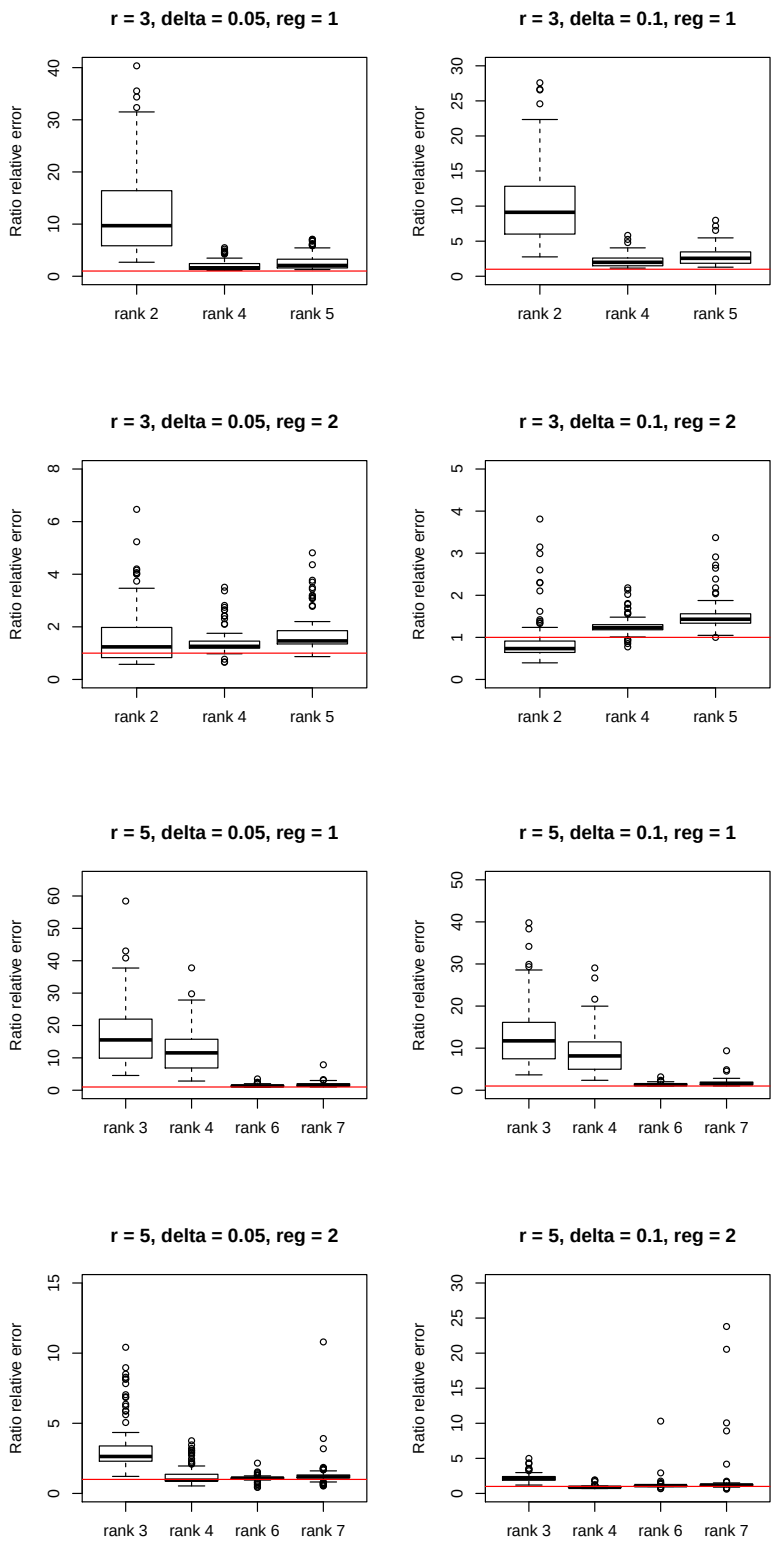
r = 3, delta = 0.05, reg = 1
r = 3, delta = 0.1, reg = 1
r = 3, delta = 0.05, reg = 2
r = 3, delta = 0.1, reg = 2
r = 5, delta = 0.05, reg = 1
r = 5, delta = 0.1, reg = 1
r = 5, delta = 0.05, reg = 2
r = 5, delta = 0.1, reg = 2
Ratio relative error
rank 2
rank 4
rank 5
rank 3
rank 6
rank 7

FIG 11. *Scenario A, combinations 3-4 and 7-8 (top four) of regime (1) and combinations 5-6 and 9-10 (bottom four) of regime (2). Underestimation is impactful in regime 1 and overestimation does not have severe impact in both regimes. Two outliers have been left out of the plots in order to allow for a better appreciation the shape of the distributions.*

Regime 1				
(n, K)	(rk,δ)	PACE	KL	RS
$n = 300, K = 25$	$(1, 0.05)$	**4.47** $(3.03, 8.60)$	**4.41** $(2.38, 5.56)$	**8.28** $(6.78, 10.2)$
	$(1, 0.10)$	**5.53** $(3.11, 9.58)$	**4.18** $(2.93, 6.54)$	**8.44** $(7.26, 11.6)$
	$(3, 0.05)$	**2.77** $(2.22, 3.67)$	**2.88** $(2.44, 3.42)$	**3.78** $(3.30, 4.49)$
	$(3, 0.10)$	**2.95** $(2.44, 3.70)$	**2.63** $(2.14, 3.17)$	**3.99** $(3.49, 4.59)$
	$(5, 0.05)$	**1.05** $(2.83, 1.32)$	**1.12** $(0.87, 1.41)$	**1.21** $(0.84, 1.56)$
	$(5, 0.10)$	0.97 $(0.61, 1.28)$	**1.00** $(0.56, 1.39)$	**1.08** $(0.72, 1.57)$
$n = 300, K = 50$	$(1, 0.05)$	**6.19** $(3.82, 9.29)$	**4.50** $(2.51, 6.41)$	**6.61** $(5.13, 7.97)$
	$(1, 0.10)$	**5.53** $(3.45, 9.10)$	**4.84** $(2.76, 6.34)$	**6.66** $(5.40, 8.16)$
	$(3, 0.05)$	**3.66** $(2.94, 5.49)$	**3.49** $(2.70, 4.22)$	**3.89** $(3.33, 4.37)$
	$(3, 0.10)$	**3.61** $(2.57, 4.89)$	**3.21** $(2.35, 4.12)$	**3.75** $(3.20, 4.45)$
	$(5, 0.05)$	**2.38** $(1.90, 3.16)$	**2.40** $(1.87, 3.02)$	**2.00** $(1.67, 2.25)$
	$(5, 0.10)$	**2.04** $(1.70, 2.80)$	**2.01** $(1.62, 2.68)$	**2.02** $(1.63, 2.47)$
$n = 300, K = 150$	$(1, 0.05)$	**3.36** $(2.04, 6.58)$	**2.62** $(1.76, 5.07)$	**3.74** $(3.00, 4.84)$
	$(1, 0.10)$	**3.19** $(1.71, 7.17)$	**2.49** $(1.41, 4.79)$	**3.85** $(2.88, 4.45)$
	$(3, 0.05)$	**2.86** $(2.23, 4.33)$	**2.63** $(2.07, 3.75)$	**2.57** $(2.22, 3.01)$
	$(3, 0.10)$	**2.76** $(1.87, 4.28)$	**2.45** $(1.73, 3.64)$	**2.41** $(1.99, 2.90)$
	$(5, 0.05)$	**2.68** $(2.00, 3.76)$	**2.47** $(1.94, 3.29)$	**1.82** $(1.59, 2.14)$
	$(5, 0.10)$	**2.08** $(1.66, 2.83)$	**1.96** $(1.53, 2.60)$	**1.58** $(1.39, 1.90)$
$n = 100, K = 25$	$(1, 0.05)$	**3.04** $(2.24, 5.02)$	**2.67** $(2.01, 3.44)$	**4.94** $(4.19, 6.25)$
	$(1, 0.10)$	**3.24** $(2.42, 4.77)$	**2.76** $(2.08, 3.44)$	**5.05** $(4.21, 6.74)$
	$(3, 0.05)$	**2.47** $(1.95, 3.02)$	**2.34** $(1.98, 2.82)$	**2.78** $(2.45, 3.15)$
	$(3, 0.10)$	**2.06** $(1.70, 2.67)$	**2.00** $(1.54, 2.42)$	**2.56** $(2.10, 2.98)$
	$(5, 0.05)$	0.74 $(0.52, 0.98)$	0.75 $(0.56, 1.02)$	0.73 $(0.48, 1.07)$
	$(5, 0.10)$	0.60 $(0.37, 0.86)$	0.62 $(0.36, 0.97)$	0.64 $(0.34, 0.98)$
$n = 100, K = 50$	$(1, 0.05)$	**3.29** $(2.21, 5.17)$	**2.70** $(2.00, 3.90)$	**3.79** $(3.19, 4.57)$
	$(1, 0.10)$	**3.08** $(1.88, 4.60)$	**2.57** $(1.59, 3.58)$	**3.85** $(2.96, 4.54)$
	$(3, 0.05)$	**2.75** $(2.11, 3.71)$	**2.70** $(2.16, 3.13)$	**2.87** $(2.54, 3.22)$
	$(3, 0.10)$	**2.25** $(1.85, 2.75)$	**2.10** $(1.78, 2.58)$	**2.42** $(2.16, 2.81)$
	$(5, 0.05)$	**1.87** $(1.56, 2.22)$	**1.91** $(1.51, 2.46)$	**1.57** $(1.28, 1.97)$
	$(5, 0.10)$	**1.68** $(1.34, 1.99)$	**1.74** $(1.35, 2.00)$	**1.49** $(1.27, 1.84)$
$n = 100, K = 100$	$(1, 0.05)$	**2.88** $(1.68, 4.43)$	**2.38** $(1.45, 3.45)$	**2.91** $(2.28, 3.79)$
	$(1, 0.10)$	**2.98** $(1.80, 4.34)$	**2.35** $(1.41, 3.37)$	**2.84** $(2.21, 3.45)$
	$(3, 0.05)$	**2.68** $(2.18, 3.52)$	**2.56** $(2.08, 3.31)$	**2.35** $(2.06, 2.85)$
	$(3, 0.10)$	**2.47** $(1.80, 3.36)$	**2.39** $(1.77, 2.97)$	**2.25** $(1.85, 2.66)$
	$(5, 0.05)$	**1.92** $(1.62, 2.58)$	**1.87** $(1.60, 2.62)$	**1.56** $(1.35, 1.79)$
	$(5, 0.10)$	**1.75** $(1.46, 2.11)$	**1.79** $(1.48, 2.21)$	**1.42** $(1.26, 1.60)$

TABLE 7

Table containing the median (the first and third quartiles are in parentheses) of the ratios for the three methods we compared our method with for different combinations of n and K with the regime 1. We highlight in bold the medians that exceed 1.

Regime 2				
(n, K)	(rk,δ)	PACE	KL	RS
$n = 300, K = 25$	$(3, 0.05)$	**1.90** $(0.64, 2.39)$	**1.49** $(0.66, 2.34)$	**2.40** $(1.16, 3.46)$
	$(3, 0.10)$	**1.64** $(1.13, 2.18)$	**1.34** $(0.99, 1.75)$	**2.46** $(1.86, 3.12)$
	$(5, 0.05)$	0.20 $(0.01, 0.83)$	0.21 $(0.01, 0.89)$	0.23 $(0.02, 0.98)$
	$(5, 0.10)$	0.19 $(0.01, 0.81)$	0.16 $(0.01, 0.72)$	0.24 $(0.01, 1.04)$
$n = 300, K = 50$	$(3, 0.05)$	**2.04** $(1.57, 2.88)$	**2.08** $(1.26, 2.92)$	**2.82** $(1.57, 3.58)$
	$(3, 0.10)$	**1.73** $(1.31, 2.07)$	**1.46** $(1.13, 2.19)$	**2.36** $(1.43, 3.07)$
	$(5, 0.05)$	**1.30** $(1.08, 1.50)$	**1.21** $(0.99, 1.59)$	**1.18** $(0.86, 1.84)$
	$(5, 0.10)$	**1.22** $(0.99, 1.41)$	**1.17** $(0.99, 1.36)$	**1.24** $(0.93, 1.71)$
$n = 300, K = 150$	$(3, 0.05)$	**1.63** $(0.95, 2.24)$	**1.77** $(0.91, 2.45)$	**1.81** $(0.85, 2.47)$
	$(3, 0.10)$	**1.02** $(0.80, 1.40)$	0.90 $(0.77, 1.43)$	**1.01** $(0.66, 1.67)$
	$(5, 0.05)$	0.87 $(0.76, 1.60)$	0.87 $(0.79, 1.63)$	0.75 $(0.42, 1.56)$
	$(5, 0.10)$	0.84 $(0.73, 1.02)$	0.83 $(0.73, 1.00)$	0.82 $(0.56, 1.20)$
$n = 100, K = 25$	$(3, 0.05)$	**1.34** $(0.90, 1.57)$	**1.12** $(0.89, 1.44)$	**1.61** $(1.24, 2.01)$
	$(3, 0.10)$	**1.25** $(0.86, 1.52)$	**1.12** $(0.76, 1.51)$	**1.77** $(1.42, 2.06)$
	$(5, 0.05)$	0.15 $(0.01, 0.80)$	0.16 $(0.01, 0.79)$	0.20 $(0.02, 0.82)$
	$(5, 0.10)$	0.55 $(0.05, 0.78)$	0.53 $(0.05, 0.81)$	0.69 $(0.06, 0.94)$
$n = 100, K = 50$	$(3, 0.05)$	**1.29** $(1.11, 1.52)$	**1.16** $(0.95, 1.64)$	**1.45** $(1.04, 2.07)$
	$(3, 0.10)$	**1.24** $(0.98, 1.62)$	**1.09** $(0.86, 1.63)$	**1.47** $(1.18, 1.91)$
	$(5, 0.05)$	**1.12** $(1.02, 1.23)$	**1.12** $(0.98, 1.23)$	**1.01** $(0.78, 1.32)$
	$(5, 0.10)$	**1.01** $(0.86, 1.17)$	**1.02** $(0.88, 1.23)$	**1.04** $(0.85, 1.26)$
$n = 100, K = 100$	$(3, 0.05)$	0.98 $(0.77, 1.49)$	0.98 $(0.75, 1.51)$	**1.07** $(0.69, 1.67)$
	$(3, 0.10)$	0.86 $(0.68, 1.13)$	0.84 $(0.64, 1.11)$	**1.90** $(0.68, 1.34)$
	$(5, 0.05)$	0.94 $(0.81, 1.09)$	0.92 $(0.81, 1.11)$	0.85 $(0.64, 1.18)$
	$(5, 0.10)$	0.88 $(0.73, 1.06)$	0.89 $(0.71, 1.11)$	0.83 $(0.59, 1.15)$

TABLE 8

Table containing the median (the first and third quartiles are in parentheses) of the ratios for the three methods we compared our method with four different combinations of n and K with the regime 2. We highlight in bold the medians that exceed 1.

Acknowledgements

We gratefully acknowledge support from Swiss National Science Foundation. Victor M. Panaretos wishes to thank Prof. Hans-Georg Müller for several stimulating discussions.

References

[1] BOSQ, D. (2000). *Linear processes in function space*. Springer.

[2] BOSQ, D. (2014). Computing the best linear predictor in a Hilbert space; applications to general ARMAH processes. *Journal of Multivariate Analysis* **124** 436–450.

[3] BOYLE, J. P. and DYKSTRA, R. L. (1986). A method for finding projections onto the intersection of convex sets in Hilbert spaces. In *Advances in order restricted statistical inference* 28–47. Springer.

[4] CHEN, Y. and WAINWRIGHT, M. J. (2015). Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees. *ArXiv e-prints*.

[5] COLEMAN, T. F. and LI, Y. (1994). On the convergence of reflective newton methods for large-scale nonlinear minimization subject to bounds. *Mathematical programming* **67** 189–224.

[6] COLEMAN, T. F. and LI, Y. (1996). An interior, trust region approach for nonlinear minimization subjects to bounds. *SIAM Journal on optimization* **6** 418–445.

[7] DAUXOIS, J., POUSSE, A. and ROMAIN, Y. (1982). Asymptotic theory for the principal component analysis of a vector random function: some applications to statistical inference. *Journal of multivariate analysis* **12** 136–154.

[8] HALL, P., MÜLLER, H.-G. and WANG, J.-L. (2006). Properties of principal component methods for functional and longitudinal data analysis. *The annals of statistics* 1493–1517.

[9] HORVÁTH, L. and KOKOSZKA, P. (2012). *Inference for functional data with applications* **200**. Springer Science & Business Media.

[10] HSING, T. and EUBANK, R. (2015). *Theoretical foundations of functional data analysis, with an introduction to linear operators*. John Wiley & Sons.

[11] Király, F. and Tomioka, R. (2012). A combinatorial algebraic approach for the identifiability of low-rank matrix completion. *Proceedings of the 29th International Conference on Machine Learning*.
[12] Krantz, S. G. and Parks, H. R. (2002). *A primer of real analytic functions*. Birkhäuser.
[13] Li, Y. and Hsing, T. (2010). Uniform convergence rates for nonparametric regression and principal component analysis in functional/longitudinal data. *The Annals of Statistics* 3321–3351.
[14] MATLAB (2012). *version 8.0.0.783 (R2012b)*. The MathWorks Inc., Natick, Massachusetts.
[15] Opsomer, J., Wang, Y. and Yang, Y. (2001). Nonparametric regression with correlated errors. *Statistical Science* 134–153.
[16] Panaretos, V. M., Kraus, D. and Maddocks, J. H. (2010). Second-order comparison of Gaussian random functions and the geometry of DNA minicircles. *Journal of the American Statistical Association* **105** 670–682.
[17] Panaretos, V. M. and Tavakoli, S. (2013). Cramér–Karhunen–Loève representation and harmonic principal component analysis of functional time series. *Stochastic Processes and their Applications* **123** 2779–2807.
[18] Ramsay, J. O. and Silverman, B. W. (2005). *Functional Data Analysis*. Springer, New York.
[19] Silverman, B. W. (1984). Spline smoothing: the equivalent variable kernel method. *The Annals of Statistics* 898–916.
[20] van der Vaart, A. W. and Wellner, J. A. (1996). *Weak convergence and empirical processes with applications to statistics*. Springer series in statistics.
[21] Wang, J. L., Chiou, J. M. and Müller, H.-G. (2015). Review of Functional Data Analysis. *ArXiv e-prints*.
[22] Yao, F., Müller, H.-G. and Wang, J.-L. (2005). Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association* **100** 577–590.
[23] Yao, F., Müller, H.-G., Wang, J.-L. et al. (2005). Functional linear regression analysis for longitudinal data. *The Annals of Statistics* **33** 2873–2903.